



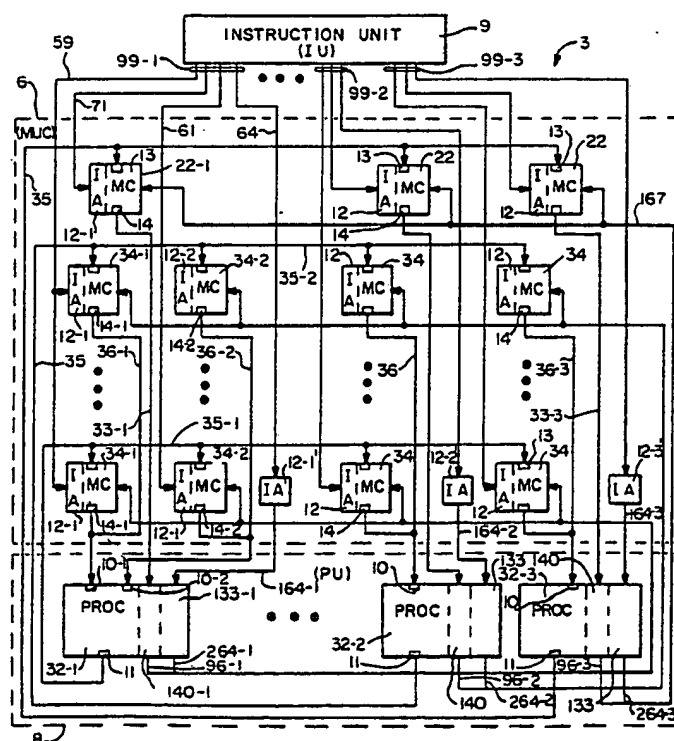
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁴ : G06F 9/38, 13/02		A1	(11) International Publication Number: WO 88/ 08568
			(43) International Publication Date: 3 November 1988 (03.11.88)
(21) International Application Number: PCT/US88/01413		(72) Inventors: RAU, Bantwal, Ramakrishna ; 159 Belvue Drive, Los Gatos, CA 95032 (US). TOWLE, Ross, Albert ; 32719 Ripon Court, Union City, CA 94587 (US). YEN, David, Wei-Luen ; 21560 Edward Way, Cupertino, CA 95014 (US). YEN, Wei-Chen ; 10431 Plum Tree Lane, Cupertino, CA 95014 (US).	
(22) International Filing Date: 30 April 1988 (30.04.88)		(74) Agents: LOVEJOY, David, E. et al.; Fliesler, Dubb, Meyer and Lovejoy, 4 Embarcadero Center, Suite 400, San Francisco, CA 94111-4156 (US).	
(31) Priority Application Numbers: 045,882 045,883 045,884 045,895 045,896		(81) Designated States: AU, DE (European patent), FR (European patent), GB (European patent), JP.	
(32) Priority Dates: 1 May 1987 (01.05.87) 1 May 1987 (01.05.87) 1 May 1987 (01.05.87) 1 May 1987 (01.05.87) 1 May 1987 (01.05.87)		Published With international search report.	
(33) Priority Country: US			
(71) Applicant: CYDRONE, INC. [US/US]; 1589 Centre Pointe Drive, Milpitas, CA 95035 (US).			

(54) Title: PARALLEL-PROCESSING SYSTEM EMPLOYING A HORIZONTAL ARCHITECTURE COMPRISING MULTIPLE PROCESSING ELEMENTS AND INTERCONNECT CIRCUIT WITH DELAY MEMORY ELEMENTS TO PROVIDE DATA PATHS BETWEEN THE PROCESSING ELEMENTS

(57) Abstract

A computer system (3) including a processing unit (8) having one or more processors (32-1, 32-2, 32-3), for performing operations on input operands and providing output operands (11), a multiconnect unit (6) for storing operands at addressable locations (34-1, 34-2) and for providing said input operands (10-1) from source addresses and for storing said output operands with destination addresses, an instruction unit (9) for specifying operations to be performed by said processing unit (8), for specifying source address offsets and destination address offsets relative to a modifiable pointer, invariant addressing means (12) for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit (6).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	ML	Mali
AU	Australia	GA	Gabon	MR	Mauritania
BB	Barbados	GB	United Kingdom	MW	Malawi
BE	Belgium	HU	Hungary	NL	Netherlands
BG	Bulgaria	IT	Italy	NO	Norway
BJ	Benin	JP	Japan	RO	Romania
BR	Brazil	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	LI	Liechtenstein	SN	Senegal
CH	Switzerland	LK	Sri Lanka	SU	Soviet Union
CM	Cameroon	LU	Luxembourg	TD	Chad
DE	Germany, Federal Republic of	MC	Monaco	TG	Togo
DK	Denmark	MG	Madagascar	US	United States of America
FI	Finland				

PARALLEL-PROCESSING SYSTEM EMPLOYING
A HORIZONTAL ARCHITECTURE COMPRISING MULTIPLE
PROCESSING ELEMENTS AND INTERCONNECT
CIRCUIT WITH DELAY MEMORY ELEMENTS TO PROVIDE DATA
PATHS BETWEEN THE PROCESSING ELEMENTS

BACKGROUND OF INVENTION

The present invention relates to computers, and more particularly, to high-speed, parallel-processing computers employing horizontal architectures.

Typical examples of computers are the IBM 360/370 Systems. In such systems, a series of general purpose registers (GPRs) are accessible to supply data to an arithmetic and logic unit (ALU). The output from the arithmetic and logic unit intern supplies results from arithmetic and logic operations to one or more of the general purpose registers. In a similar manner, some 360/370 Systems include a floating point processor (FPP) and include corresponding floating point registers (FPRs). The floating point registers supply data to the floating point processor and, similarly, the results from the floating point processor are stored back into one or more of the floating point registers. The types of instructions which employ either the GPRs or the FPRs are register to register (RR) instructions. Frequently, in the operation of the GPRs and the FPRs for RR instructions, identical data is stored in two or more register locations. Accordingly, the operation storing into the GPRs is selectively to one or more locations. Similarly, the input to the ALU frequently is selectively from one or more of many locations storing the same data.

Horizontal processors have been proposed for a number of years. See for example, "SOME SCHEDULING TECHNIQUES AND

-2-

AN EASILY SCHEDULABLE HORIZONTAL ARCHITECTURE FOR HIGH PERFORMANCE SCIENTIFIC COMPUTING" by B.R. Rau and C.D. Glaeser, IEEE Proceedings of the 14th Annual Microprogramming Workshop, October 1981, pp. 1983-198 Advanced Processor Technology Group ESL, Inc., San Jose, California, and "EFFICIENT CODE GENERATION FOR HORIZONTAL ARCHITECTURES: COMPILER TECHNIQUES AND ARCHITECTURAL SUPPORT" BY B. Ramakrishna Rau, Christopher D. Glaeser and Raymond L. Picard, IEEE 9th Annual Symposium on Computer Architecture 1982, pp. 131-139.

Horizontal architectures have been developed to perform high speed scientific computations at a relatively modest cost. The simultaneous requirements of high performance and low cost lead to an architecture consisting of multiple pipelined processing elements (PEs), such as adders and multipliers, a memory (which for scheduling purposes may be viewed as yet another PE with two operations, a READ and WRITE), and an interconnect which ties them all together.

The interconnect allows the result of one operation to be directly routed to one of the inputs for another processing elements where another operation is to be performed. With such an interconnect the required memory bandwidth is reduced since temporary values need not be written to and read from the memory. Another aspect typical of horizontal processors is that their program memories emit wide instructions which synchronously specify the actions of the multiple and possibly dissimilar processing elements. The program memory is sequenced by a sequencer that assumes sequential flow of control unless a branch is explicitly specified.

As a consequence of their simplicity, horizontal architectures are inexpensive when considering the potential performance obtainable. However, if this potential performance is to be realized, the multiple resources of a horizontal processor must be scheduled effectively. The

-3-

scheduling task for conventional horizontal processors is quite complex and the construction of highly optimizing compilers for them is difficult and expensive.

The polycyclic architecture has been designed to support code generation by simplifying the task of scheduling the resources of horizontal processors. The advantages are 1) that the scheduler portion of the compiler will be easier to implement, 2) that the code generated will be of a higher quality, 3) that the compiler will execute fast, and 4) that the automatic generation of compilers will be facilitated.

The polycyclic architecture is a horizontal architecture that has unique interconnect and delay elements. The interconnect element of a polycyclic processor has a dedicated delay element between every directly connected resource output and resource input. This delay element enables a datum to be delayed by an arbitrary amount of time in transit between the corresponding output and input.

The topology of the interconnect may be arbitrary. It is possible to design polycyclic processors with n resources in which the number of delay elements is $O(n)$, (a uni- or multi-bus structure), $O(n \log n)$, (e.g. delta networks, or $O(n^2)$, (a cross-bar). The trade-offs are between cost, interconnect bandwidth and interconnect latency. Thus, it is possible to design polycyclic processors lying in various cost-performance brackets.

In previously proposed polycyclic processors, the structure of an individual delay element consists of a register file, any location of which may be read by providing an explicit read address. Optionally, the value accessed can be deleted. This option is exercised on the last access to that value. The result is that every value with addresses greater than the address of the deleted value is simultaneously shifted down, in one machine cycle, to the location with the next lower address. Consequently, all values present in the delay element are compacted into the

-4-

lowest locations. An incoming value is written into the lowest empty location which is always pointed to by the Write Pointer that is maintained by hardware. The Write Pointer is automatically incremented each time a value is written and is decremented each time one is deleted. As a consequence of deletions, a value, during its residence in the delay element, drifts down to lower addresses, and is read from various locations before it is itself deleted.

A value's current position at each instant during execution must be known by the compiler so that the appropriate read address may be specified by the program when the value is to be read. Keeping track of this position is a tedious task which must be performed by a compiler during code-generation.

To illustrate the differences, two processors, a conventional horizontal processor and a polycyclic processor, are compared. A typical horizontal processor contains one adder and one multiplier, each with a pipeline stage time of one cycle and a latency of two cycles. It also contains two scratch-pad register files labeled A and B. The interconnect is assumed to consist of a delayless cross-bar with broadcast capabilities, that is, the value at any input port may be distributed to any number of the output ports simultaneously. Each scratch-pad is assumed to be capable of one read and one write per cycle. A read specified on one cycle causes the datum to be available at the output ports of the interconnect on the next cycle. If a read and a write with the same address are specified on the same scratch-pad on the same cycle, then the datum at the input of the scratch-pad during that cycle will be available at the output ports of the interconnect on the next cycle. In this manner, a delay of one or more cycles may be obtained in transmitting a value between the output of one processor and the input of another. The horizontal processor typically also contains other resources.

-5-

A typical polycyclic processor is similar to the horizontal processor except for the nature of the interconnect element and the absence of the two scratchpad register files. While the horizontal processor's interconnect is a crossbar, the polycyclic processor's interconnect is a crossbar with a delay element at each cross-point. The interconnect has two output ports (columns) and one input port (row) for each of the two processing elements. Each cross-point has a delay element which is capable of one read and one write each cycle.

In previously proposed processors, a processor can simultaneously distribute its output to any or all of the delay elements which are in the row of the interconnect corresponding to its output port. A processor can obtain its input directly. If a value is written into a delay element at the same time that an attempt is made to read from the delay element, the value is transmitted through the interconnect with no delay. Any delay may be obtained merely by leaving the datum in the delay element for a suitable length of time.

In the polycyclic processors proposed, elaborate controls were provided for determining which delay element in a row actually received data as a result of a processor operation and the shifting of data from element to element. This selection process and shifting causes the elements in a row to have different data at different times. Furthermore, the removal of data from the different delay elements requires an elaborate process for purging the data elements at appropriate times. The operations of selecting and purging data from the data elements is somewhat cumbersome and is not entirely satisfactory.

Although the polycyclic and horizontal processors which have previously been proposed are an improvement over previously known systems, there still is a need for still additional improvements which increase the performance and efficiency of processors.

-6-

SUMMARY OF INVENTION

The present invention is a horizontal architecture computer system including a processing unit, a multiconnect unit (row and column register file), an instruction unit, and an invariant addressing unit.

The processing unit performs operations on input operands and provides output operands. The multiconnect unit stores operands at multiconnect address locations and provides the input operands to the processing unit from multiconnect source addresses and stores the output operands from the processing unit at multiconnect destination addresses. The instruction unit specifies operations to be performed by the processing unit and specifies multiconnect address offsets for the operands in the multiconnect unit relative to multiconnect pointers. The invariant addressing unit combines pointers and address offsets to form the actual addresses of operands in the multiconnect unit. The pointer is modified to to sequence the actual address locations accessed in the multiconnect unit.

Typically, the processing unit includes conventional processors such as adders, multipliers, memories and other functional units for performing operations on operands and for storing and retrieving operands under program control.

The multiconnect is a register file formed of addressable memory circuits forming multiconnect elements organized in rows and columns. Each multiconnect element (memory circuit) has addressable multiconnect locations for storing operands.

The multiconnect elements are organized in columns such that a column of multiconnect elements are connected to a common data bus providing a data input to a processor. Each multiconnect element in a column, when addressed, provides a source operand to the common data bus in response to a source address. Each processor receives one

-7-

or more data input buses from a column of multiconnect elements. Each column of multiconnect elements is addressed by a different source address formed by combining a different source offset from the instruction unit with the multiconnect pointer.

The multiconnect elements are organized in rows such that each processor has an output connected as an input to a row of multiconnect elements so that processor output operands are stored identically in each multiconnect element in a row. The particular location in which an output operand is stored in each multiconnect element in a row is specified by a destination address formed by combining a destination offset from the instruction unit with the multiconnect pointer.

Each address offset, including the column source address offset and the row destination address offset, is specified relative to a multiconnect pointer (mcp).

The invariant addressing unit combines the instruction specified address offset with the multiconnect pointer to provide the actual address of the source or destination operand in the multiconnect unit.

The multiconnect unit permits each processor to receive operands as an input from the output of any other processor at the same time. The multiconnect permits the changing of actual address locations by changing the pointer without changing the relative location of operands.

The instructions executed by the computer of the present invention are scheduled to make efficient use of the available processors and other resources in the system and to insure that no conflict exists for use of the resources of the system.

As a result of scheduling a program, an initial instruction stream, IS, of scheduled instructions is formed. Each initial instruction, I_i , in the initial instruction stream is formed by a set of zero, one or more operations that are to be initiated concurrently. When an

-8-

instruction specifies only a single operation, it is a single-operation instruction and when it specifies multiple operations, it is a multi-operation instruction.

The initial instructions in the initial instruction stream, IS, are transformed to a transformed (or kernel) instruction stream, \bar{IS} , having Y transformed (kernel) instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \bar{I}_3, \dots, \bar{I}_k, \dots, \bar{I}_{(Y-1)}$ where $0 \leq k \leq (Y-1)$.

Each kernel instruction, \bar{I}_k , in the kernel instruction stream \bar{IS} is formed by a set of zero, one or more operations, $\bar{O}_0^{k,l}, \bar{O}_1^{k,l}, \bar{O}_2^{k,l}, \dots, \bar{O}_n^{k,l}, \dots, \bar{O}_{(N-1)}^{k,l}$, initiated concurrently where $0 \leq n \leq (N-1)$, where N is the number of processors for performing operations and where the kernel operation, $\bar{O}_n^{k,l}$, is performed by the n^{th} -processor in response to the k^{th} -kernel instructions.

An initial instruction stream, IS, is frequently of the type having a loop, LP, in which the L instructions forming the loop are repeatedly executed a number of times, R, during the processing of the instruction stream. After transformation, an initial loop, LP, is converted to a kernel loop, KE, of K kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ in which execution sequences from \bar{I}_0 toward $\bar{I}_{(K-1)}$ one or more times, once for each execution of the kernel, KE.

The computer system executes a loop with overlapped code. Iteration control circuitry is provided for selectively controlling the operations of the kernel instructions. Different operations specified by each kernel instruction are initiated as a function of the particular iteration of the loop that is being performed. The iterations are partitioned into a prolog, body, and epilog. During successive prolog iterations, an increasing number of operations are performed, during successive body iterations, a constant number of operations are performed, and during successive epilog iterations, a decreasing number of operations are performed.

-9-

The iteration control circuitry includes controls for counting the iterations of a loop, the prolog iterations, the body iterations and the epilog iterations. In one particular embodiment, a loop counter counts the loops and an epilog counter counts the iterations during the epilog. An iteration control register is provided for controlling each processor to determine which operations are active during each iteration.

The computer system efficiently executes loops of instructions with recurrence, that is, where the results from one iteration of the loop are used in subsequent iterations of the loop.

The iteration control circuitry includes controls for counting the iterations of a loop, the prolog iterations, the body iterations and the epilog iterations. In one particular embodiment, a loop counter counts the loops and an epilog counter counts the iterations during the epilog. An iteration control register is provided for controlling each processor to determine which operations are active during each iteration.

The computer system efficiently executes loops of instructions with branch in the loop, that is, where the instruction path from one iteration of the loop may be different in subsequent iterations of the loop.

The iteration control circuitry includes controls for counting the iterations of a loop, the prolog iterations, the body iterations and the epilog iterations. In one particular embodiment, a loop counter counts the loops and an epilog counter counts the iterations during the epilog. An iteration control register is provided for controlling each processor to determine which operations are active during each iteration.

The computer system executes a loop with overlapped code. The instruction unit has a plurality of locations for specifying an operation to be performed by said processors, each for specifying source address offsets and

-10-

destination address offsets relative to a modifiable pointer.

BRIEF DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of an overall system in accordance with the present invention.

FIG. 2 is a block diagram of a numeric processor computer which forms part of the FIG. 1 system.

FIG. 3 is a block diagram of one preferred embodiment of a numeric processor computer of the FIG. 2 type.

FIG. 4 is a block diagram of the instruction unit which forms part of the computer of FIG. 3.

FIG. 5 is a block diagram of the instruction address generator which forms part of the instruction unit of FIG. 4.

FIG. 6 is a block diagram of an invariant addressing unit which is utilized within the computer of FIGS. 2 and 3.

FIG. 7 depicts a schematic representation of a typical processor of the type employed within the processing unit of FIG. 2 or FIG. 3.

FIG. 8 depicts the STUFFICR control processor for controlling predicate values within the FIG. 3 system.

FIG. 9 depicts a multiply processor within the FIG. 3 system.

FIG. 10 is a block diagram of typical portion of the mutliconnect and the corresponding processors which form part of the system of FIG. 3.

FIG. 11 depicts a block diagram of one preferred implementation of a physical mutliconnect which forms one half of two logical mutliconnects within the FIG. 3 system.

FIGS. 12 and 13 depict electrical block diagrams of portions of the physical multiconnect of FIG. 11.

-11-

FIG. 14 depicts a block diagram of a typical predicate ICR multiconnect within the FIG. 3 system.

FIG. 15 depicts a block diagram of an instruction unit for use with multiple operation and single operation instructions.

DETAILED DESCRIPTION

Overall System - FIG. 1

In FIG. 1, a high performance, low-cost system 2 is shown for computation-intensive numeric tasks. The FIG. 1 system processes computation tasks in the numeric processor(NP) computer 3.

The computer 3 typically includes a processing unit(PU) 8 for computation-intensive task, includes an instruction unit(IU) 9 for the fetching, dispatching, and caching of instructions, includes a register multiconnect unit (MCU) 6 for connecting data from and to the processing unit 8, and includes an interface unit(IFU) 23 for passing data to and from the main store 7 over bus 5 and to and from the I/O 24 over bus 4. In one embodiment, the interface unit 23 is capable of issuing two main store requests per clock for the multiconnect unit 6 and one request per clock for the instruction unit 9.

The computer 3 employs a horizontal architecture for executing an instruction stream, \bar{IS} , fetched by the I unit 9. The instruction stream includes a number of instructions, $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ where each instruction, \bar{I}_k , of the instruction stream \bar{IS} specifies one or more operations $\bar{O}_1^{k,l}, \bar{O}_2^{k,l}, \dots, \bar{O}_n^{k,l}, \dots, \bar{O}_N^{k,l}$, to be performed by the processing unit 8.

In one embodiment, the processing unit 8 includes a number, N, of parallel processors, where each processor performs one or more of the operations, $\bar{O}_n^{k,l}$. Each instruction from the instruction unit 9 provides source addresses (or source address offsets) for specifying the addresses of

-12-

operands in the multiconnect unit 6 to be transferred to the processing unit 8. Each instruction from the instruction unit 9 provides destination addresses (or destination address offsets) for specifying the addresses in the multiconnect unit 6 to which result operands from the processing unit 8 are to be transferred. The multiconnect unit 6 is a register file where the registers are organized in rows and columns and where the registers are accessed for writing into in rows and are accessed for reading from in columns. The columns connect information from the multiconnect unit 6 to the processing unit 8 and the rows connect information from the processing unit to the multiconnect 6.

The source and destination addresses of the multiconnect from the instruction unit 9 are specified using invariant addressing. The invariant addressing is carried out in invariant addressing units (IAU) 12 which store multiconnect pointers(mcp). The instructions provide addresses in the form of address offsets(ao) and the address offsets(ao) are combined with the multiconnect pointers(mcp) to form the actual source and destination addresses in the multiconnect. The locations in multiconnect unit are specified with mcp-relative addresses.

Instruction Scheduling

The execution of instructions by the computer of the present invention requires that the instructions of a program be scheduled, for example, by a compiler which compiles prior to execution time. The object of the scheduling is to make efficient use of the available processors and other resources in the system and to insure that no conflict exists for use of the resources of the system. In general, each functional unit (processor) can be requested to perform only a single operation per cycle and each bus can be requested to make a single transfer per

-13-

cycle. Scheduling attempts to use, at the same time, as many of the resources of the system as possible without creating conflicts so that execution will be performed in as short a time as possible.

As a result of scheduling a program, an initial instruction stream, IS, of scheduled instructions is formed and is defined as the Z initial instructions $I_0, I_1, I_2, I_3, \dots, I_\ell, \dots, I_{(Z-1)}$ where $0 \leq \ell \leq Z$. The scheduling to form the initial instruction stream, IS, can be performed using any well-known scheduling algorithm. For example, some methods for scheduling instructions are described in the publications listed in the above BACKGROUND OF INVENTION.

Each initial instruction, I_ℓ , in the initial instruction stream is formed by a set of zero, one or more operations, $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, that are to be initiated concurrently, where $0 \leq n \leq (N-1)$, where N is the number of processors for performing operations and where the operation O_n^ℓ is performed by the n^{th} -processor in response to the ℓ^{th} -initial instruction. When an instruction has zero operations, the instruction is termed a "NO OP" and performs no operations. When an instruction specifies only a single operation, it is a single-op instruction and when it specifies multiple operations, it is a multi-op instruction.

In accordance with the present invention, the initial instructions in the initial instruction stream, IS, are transformed to a transformed(or kernel) instruction stream, \bar{IS} , having Y transformed(kernel) instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \bar{I}_3, \dots, \bar{I}_k, \dots, \bar{I}_{(Y-1)}$ where $0 \leq k \leq (Y-1)$.

Each kernel instruction, \bar{I}_k , in the kernel instruction stream \bar{IS} is formed by a set of zero, one or more operations, $\bar{O}_0^{k,\ell}, \bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_{(N-1)}^{k,\ell}$, initiated concurrently where $0 \leq n \leq (N-1)$, where N is the number of processors for performing operations and where the kernel operation, $\bar{O}_n^{k,\ell}$, is performed by the n^{th} -processor in response to the k^{th} -kernel instructions.

-14-

The operations designated as $\bar{O}_0^{k,\ell}, \bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_{(N-1)}^{k,\ell}$ for the kernel k^{th} -instruction, \bar{I}_k , correspond to selected ones of the operations $O_0^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$ selected from all L of the initial instructions I_ℓ for which the index k satisfies the following:

$$k = \ell \text{MOD}[K]$$

where, $0 \leq \ell \leq (L-1)$.

Each kernel operation $\bar{O}_n^{k,\ell}$ is identical to a unique one of the initial operations O_n^ℓ where the value of ℓ is given by $k = \ell \text{MOD}[K]$.

An initial instruction stream, IS , is frequently of the type having a loop, LP , in which the L instructions forming the loop are repeatedly executed a number of times, R , during the processing of the instruction stream. After transformation, an initial loop, LP , is converted to an kernel loop, KE , of K kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ in which execution sequences from \bar{I}_0 toward $\bar{I}_{(K-1)}$ one or more times, once for each execution of the kernel, KE , where $0 \leq k \leq (K-1)$.

Overlapping of Loops

A modulo scheduling algorithm, for example as described in the articles referenced under the BACKGROUND OF INVENTION, is applied to a program loop that consists of one basic block. The operations of each iteration are divided into groups based on which stage, S_j , they are in where j equals $0, 1, 2, \dots$, and so on. Thus, operations scheduled for the first iteration interval (II) of instruction cycles are in the first stage (S_0), those scheduled for the next II cycles are in the second stage (S_1), and so on. The modulo scheduling algorithm assigns the operations, O_n^ℓ , to the various stages and schedules them in such a manner that all the stages, one each from consecutive iterations, can be executed simultaneously.

-15-

As an example, consider a schedule consisting of three stages S_0 , S_1 , and S_2 per iteration. These three stages are executed in successive iterations ($i=0, 1, 2, 3, 4$) in the following example:

Iteration #	0	1	2	3	4
II_0	$S_0(0)$				
II_1	$S_1(0)$	$S_0(1)$			
II_2	$S_2(0)$	$S_1(1)$	$S_0(2)$		
II_3		$S_2(1)$	$S_1(2)$	$S_0(3)$	
II_4			$S_2(2)$	$S_1(3)$	$S_0(4)$
:				:	:
:				:	:

$S_j(i)$ represents all the operations in the j -th stage of iteration i . Starting with the interval in which the last stage of the first iteration is executed (II_2 in the above example) a repetitive execution pattern exists until the interval in which the first stage of the last iteration is executed (II_4 in the above example). All these repetitive steps may be executed by iterating with the computation of the form:

$$S_2(i-2) \ S_1(i-1) \ S_0(i)$$

This repeated step is called the kernel. The steps preceding the kernel (II_0 , II_1 in the example) are called the prolog, and the steps following the kernel (after II_4 in the example) are called the epilog. The nature of the overall computation is partitioned as follows:

-16-

```
Prolog:      S0(0)
              S1(0)  S0(1)

Kernel:      S2(i-2) S1(i-1) S0(i) for i = 3, ..., n

Epilog:      S2(n-1) S1(n)
              S2(n)
```

To avoid confusion between the iterations of the initial loop, LP, and the iterations of the kernel loop, KE, the terms ℓ -iteration and k-iteration are employed, respectively.

Assuming that the computation specified by each ℓ -iteration is identical, the modulo scheduling algorithm guarantees (by construction) that corresponding stages in any two iterations will have the same set of operations and the same relative schedules. However, if the body of the loop contains conditional branches, successive iterations can perform different computations. These different computations are handled in the kernel-only code by means of a conditionally executed operation capability which produces the same result as the situation in which each ℓ -iteration performs the same computation.

Each operation that generates a result must specify the destination address of the multiconnect(register) in which the result is to be held. In general, due to the presence of recurrences and the overlapping of ℓ -iterations, the result generated by the operation in one ℓ -iteration may not have been used (and hence still must be saved) before the result of the corresponding operation in the next ℓ -iteration is generated. Consequently, these two results cannot be placed in the same register which, in turn, means that corresponding operations in successive iterations will have different destination (and hence, source) fields. This problem, if not accounted for, prevents the overlapping of ℓ -iterations.

-17-

One way to handle this problem is to copy the result of one ℓ -iteration into some other register before the result for the next ℓ -iteration is generated. The obvious disadvantage of such copying is that a large number of copy operations result and performance is degraded. Another way around the problem without loss of performance, but with expensive hardware cost, is to organize the registers as a shift register with random read-write capability. If the registers are shifted every time a new ℓ -iteration is started, the necessary copy operations will be performed simultaneously for all results of previous ℓ -iterations. The disadvantage is the cost.

In accordance with the present invention, the problem is avoided less expensively by using register files (called multiconnects) with the source and destination fields as relative displacements from a base multiconnect pointer (mcp) when computing the read and write addresses for the multiconnect. The base pointer is modified (decremented) each time an iteration of the kernel loop is started.

Since the source and destination register addresses are determined by adding the corresponding operation address offset (ao) to the current value of the multiconnect pointer (mcp), the specification of a source address must be based on a knowledge of the destination address of the operation that generated the relevant result and the number of times that the base multiconnect pointer has been decremented since the result generating operation was executed.

Kernel-Only Code

Code space can be saved relative to the non-overlapping version of code by using overlapping kernel-only schedules. The use of kernel-only code requires the following conditions:

1. that $S_j(i)$ be identical to $S_j(k)$ for all i and k and

-18-

2. the capability of conditionally executing operations

The notation

$$S_j(i) \text{ if } p(i)$$

means that the operations in the j -th stage of the i -th iteration are executed normally if and only if the predicate (Boolean value) $p(i)$ is true. If $p(i)$ is false, every operation in $S_j(i)$ will effectively be a NOOP, that is, no operation.

These conditions yield the following kernel-only loop:

LOOP: (for $i = 1, \dots, n+2$)

$S_2(i-2), 0, 1, 2$ if $p(i-2)$

$S_1(i-1), 0, 1$ if $p(i-1)$

$S_0(i), 0$ if $p(i)$

If i is less than n

then decrement mcp pointer,

set $p(i+1)$ to true and goto LOOP

else decrement base pointer,

set $p(i+1)$ to false and goto LOOP

The initial conditions for the kernel-only code are that all the $p(i)$ have been cleared to false and $p(1)$ is set to true. After $n+2$ iterations, the loop is exited.

To provide consistency, the terms used in this specification are generally as set forth in the following term table.

Term Table

-19-

- IS = initial instruction stream having Z initial instructions $I_0, I_1, I_2, I_3, \dots, I_\ell, \dots, I_{(Z-1)}$ where $0 \leq \ell < Z$.
- I_ℓ = ℓ^{th} -initial instruction in the initial instruction stream IS formed by a set of zero, one or more operations, $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, initiated concurrently, where $0 \leq n \leq (N-1)$, where N is the number of processors for performing operations and where the operation O_n^ℓ is performed by the n^{th} -processor in response to the ℓ^{th} initial instruction. When an instruction has zero operations, the instruction is termed a "NO OP" and performs no operations.
- LP = an initial loop of L initial instructions $I_0, I_1, I_2, \dots, I_\ell, \dots, I_{(L-1)}$ which forms part of the initial instruction stream IS and in which execution sequences from I_0 toward $I_{(L-1)}$, and which commences with I_0 one or more times, once for each iteration of the loop, LP, where $0 \leq \ell \leq (L-1)$.
- L = number of instructions in the initial loop, LP.
- O_n^ℓ = n^{th} -operation of a set of zero, one or more operations $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$ for the ℓ^{th} -initial instruction, I_ℓ , where $0 \leq n \leq (N-1)$ and where N is the number of processors for performing operations and where the operation O_n^ℓ is performed by the n^{th} -processor in response to the ℓ^{th} -initial instruction.
- $\bar{I}S$ = kernel instruction stream having Y kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \bar{I}_3, \dots, \bar{I}_k, \dots, \bar{I}_{(Y-1)}$ where $0 \leq k \leq (Y-1)$.

-20-

\bar{I}_k = k^{th} -kernel instruction in the kernel instruction stream $\bar{I}S$ formed by a set of zero, one or more operations, $\bar{O}_0^{k,\ell}$, $\bar{O}_1^{k,\ell}$, $\bar{O}_2^{k,\ell}$, ..., $\bar{O}_n^{k,\ell}$, ..., $\bar{O}_{(N-1)}^{k,\ell}$, initiated concurrently where $0 \leq n \leq (N-1)$, where N is the number of processors for performing operations and where the kernel operation, $\bar{O}_n^{k,\ell}$, is performed by the n^{th} -processor in response to the k^{th} -kernel instructions. When an instruction has zero operations, the instruction is termed a "NO. OP." and performs no operations.

KE = a kernel loop of K kernel instructions \bar{I}_0 , \bar{I}_1 , \bar{I}_2 , ..., \bar{I}_k , ..., $\bar{I}_{(K-1)}$ in which execution sequences from \bar{I}_0 toward $\bar{I}_{(K-1)}$ one or more times, once for each execution of the kernel, KE , where $0 \leq k \leq (K-1)$.

K = number of instructions in kernel loop, KE .

$\bar{O}_n^{k,\ell}$ = n^{th} -operation of a set of zero, one or more operations $\bar{O}_0^{k,\ell}$, $\bar{O}_1^{k,\ell}$, $\bar{O}_2^{k,\ell}$, ..., $\bar{O}_n^{k,\ell}$, ..., $\bar{O}_{(N-1)}^{k,\ell}$ for the k^{th} -instruction, \bar{I}_k , where $0 \leq n \leq (N-1)$ and N is the number of processors for performing operations.

The operations designated as $\bar{O}_0^{k,\ell}$, $\bar{O}_1^{k,\ell}$, $\bar{O}_2^{k,\ell}$, ..., $\bar{O}_n^{k,\ell}$, ..., $\bar{O}_{(N-1)}^{k,\ell}$ for the kernel k^{th} -instruction, \bar{I}_k , correspond to selected ones of the operations O_0^ℓ , O_2^ℓ , ..., O_n^ℓ , ..., $O_{(N-1)}^\ell$ selected from all L of the initial instructions I_ℓ for which the index k satisfies the following:

$$k = \ell \text{MOD}[K]$$

where, $0 \leq \ell \leq (L-1)$

-21-

Each kernel operation $O_n^{k,\ell}$ is identical to a unique one of the initial operations O_n^ℓ where the value of ℓ is given by $k = \ell \text{MOD}[K]$.

S_n^ℓ = stage number of each n^{th} operation, $O_n^{k,\ell}$, for $0 \leq n \leq (N-1)$ in the ℓ^{th} -initial instruction, I_ℓ .

= $\text{INT}[\ell/K]$

where, $0 \leq S_n^\ell \leq (J-1)$
 $(J-1) = \text{INT}[L/K]$.

= $po_n^{k,\ell}$

J = the number of stages in the initial loop, LP.

i = iteration number for initial loop, LP.

\bar{i} = iteration number for kernel loop, KE.

II = iteration interval formed by a number, K , of instruction periods.

T = sequence number indicating cycles of the computer clock

$icp(\bar{i})$ = iteration control pointer value during \bar{i}^{th} -iteration.

$mcp(\bar{i})$ = multiconnect pointer value during \bar{i}^{th} -iteration

= $\begin{cases} \text{constant, for } \bar{i}=1 \\ D^*[mcp(\bar{i}-1)], \text{ for } \bar{i} \text{ greater than } 1. \end{cases}$

-22-

$D^*[]$ = operator for forming a modified value of $mcp(\bar{i})$ or $icp(\bar{i})$ from the previous value $mcp(\bar{i}-1)$ or $icp(\bar{i}-1)$.

$ao_n^k(c)$ = address offset for c^{th} -connector port specified by n^{th} -operation of the kernel k^{th} -instruction, \bar{i}_k .

$a_n^k(c)(\bar{i})$ = multiconnect memory address for c^{th} -connector port determined for n^{th} -operation of k^{th} -instruction during \bar{i}^{th} -iteration.

$$= ao_n^k(c) + mcp(\bar{i}).$$

$po_n^{k,\ell}$ = predicate offset specified by n^{th} -operation, $\bar{O}_n^{k,\ell}$, of kernel k^{th} -instruction.

$$= INT[\ell/K], \text{ where } 0 \leq po_n^{k,\ell} \leq (J-1).$$

The predicate offset, $po_n^{k,\ell}$, from the kernel operation $O_n^{k,\ell}$ is identical to the stage number S_n^ℓ from the initial operation. O_n^ℓ , which corresponds to the kernel operation, $O_n^{k,\ell}$. The operation O_n^ℓ corresponds to $O_n^{k,\ell}$ when for both operations ℓ equals ℓ and n equals n but k may or may not equal k .

$p_n^k(\bar{i})$ = iteration control register (icr) multiconnect memory address determined for n^{th} -operation of k^{th} -instruction during \bar{i} -iteration.

$$= po_n^{k,\ell} + icp(\bar{i}).$$

$O_n^\ell[i]$ = execution of O_n^ℓ , during i^{th} -iteration where O_n^ℓ is the n^{th} -operation within ℓ^{th} -initial instruction.

$\bar{O}_n^{k,\ell}[\bar{i}]$ = execution of $O_n^{k,\ell}$ during \bar{i}^{th} -iteration where $O_n^{k,\ell}$ is the n^{th} -operation within k^{th} -kernel

-23-

instruction.

$$\bar{i} = i + S_n^\ell$$

$$i = \bar{i} - p_n^{k, \ell}$$

$$k = \ell \text{MOD}[K]$$

lc = loop count for counting each iteration of kernel loop, KE, corresponding to iterations of initial loop, LP.

$$= (R1)$$

esc = epilog stage count for counting additional iterations of kernel loop, KE, after iterations which correspond to initial loop, LP.

$$= S_n^\ell \text{ for } \ell=L \text{ (the largest stage number).}$$

psc = prolog stage count for counting first $(S_n^\ell - 1)$ iterations of initial loop, LP.

$C_n^k(\bar{i})$ = iteration control value for n^{th} -operation during \bar{i}^{th} -iteration accessed from a control register at $p_n^k(\bar{i})$ address.

R = number of iterations of initial loop, LP, to be performed.

\bar{R} = number of iterations of kernel loop, KE to be performed.

$$\bar{R} = R + \text{esc}$$

Numeric Processor - FIG. 2

A block diagram of the numeric processor, computer 3, is shown in FIG. 2. The computer 3 employs a horizontal architecture for use in executing an instruction stream fetched by the instruction unit 9. The instruction stream includes a number of kernel instructions, $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ of an instruction stream, $\bar{I}S$, where each said instruction, \bar{I}_k , of the instruction stream $\bar{I}S$ specifies one or more operations $\bar{O}_1^{k,l}, \bar{O}_2^{k,l}, \dots, \bar{O}_n^{k,l}, \dots, \bar{O}_N^{k,l}$, where each operation, $\bar{O}_n^{k,l}$, provides address offsets, $a_n^k(c)$, used in the invariant address (IA) units 12.

To process instructions, the instruction unit 9 sequentially accesses the kernel instructions, \bar{I}_k , and corresponding operations, $\bar{O}_n^{k,l}$, one or more times during one or more iterations, \bar{I} , of the instruction stream $\bar{I}S$.

The computer 3 includes one or more processors 32, each processor for performing one or more of operations, $O_n^{k,l}$, specified by the instructions, \bar{I}_k , from the I unit 9. The processors 32 include input ports 10 and output ports 11.

The computer 3 includes a plurality of multiconnects(registers) 22 and 34, addressed by memory addresses, $a_n^k(c)(\bar{I})$, from invariant addressing(IA) units 12. The multiconnects 22 and 34 connect operands from and to the processors 32. The multiconnects 32 have input ports 13 and output ports 14. The multiconnects 34 provide input operands for the processors 32 on the memory output ports 14 when addressed by invariant addresses from the IA units 12.

The computer 3 includes processor-multiconnect buses 35 for connecting output result operands from processor output ports 11 to memory input ports 13.

The computer 3 includes multiconnect-processor buses 36 for connecting input operands from multiconnect output ports 14 to processor input ports 10.

-25-

The computer 3 includes an invariant addressing(IA) unit 12 for addressing the multiconnects 34 during different iterations including a current iteration, \bar{i} , and a previous iteration, $(\bar{i}-1)$.

In FIG. 2, the output 99-1 lines from the instruction unit 9 are associated with the processor 32-1. The S1 source address on bus 59 addresses through an invariant address unit 12 a first column of multiconnects to provide a first operand input on bus 36-1 to processor 32-1 and the S2 source address on bus 61 addresses through an invariant address unit 12 a second column of multiconnects to provide a second operand input to processor 32-1 on column bus 36-2. The D1 destination address on bus 64 connects through the invariant address unit 12-1 and latency delay 133-1 to address the row of multiconnects 34 which receive the result operand from processor 32-1. The instruction unit 9 provides a predicate address on bus 71 to a predicate multiconnect 22-1 which in response provides a predicate operand on bus 33-1 to the predicate control 140-1 of processor 32-1.

In a similar manner, processors 32-2 and 32-3 of FIG. 2 have outputs 99-2 and 99-3 for addressing through invariant address units 12 the rows and columns of the multiconnect unit 6. The outputs 99-3 and processor 32-3 is associated with the multiconnect units 22 which, in one embodiment, function as the predicate control store. Processor 32-3 is dedicated to controlling the storing of predicate control values to the multiconnect 22. These control values enable the computer of FIG. 2 to execute kernel-only code, to process recurrences on loops and to process conditional recurrences on loops.

NP Processing Unit - FIG. 3

In FIG. 3, further details of the computer 3 of FIG. 2 are shown. In FIG. 3, a number of processors 32-1 through 32-9 forming the processing unit 8 are shown. The

-26-

processors 32-1 through 32-4 form a data cluster for processing data.

In the data cluster, the processor 32-1 performs floating point (FAdd) adds and arithmetic and logic unit (ALU) operations such as OR, AND, and compares including "greater than" (Fgt), "greater than or equal to" (Fge), and "less than" (Flt), on 32-bit input operands on input buses 36-1 and 37-1. The processor 32-1 forms a 32-bit result on the output bus 35-1. The bus 35-1 connects to the general purpose register (GPR) input bus 65 and connects to the row 237-1 (dmc 1) of multiconnects 34. The processor 32-1 also receives a predicate input line 33-1 from the predicate multiconnect ICR(1) in the ICR predicate multiconnect 29.

The processor 32-2 functions to perform floating point multiplies (FMpy), divides (FDiv) and square roots (FSqr). Processor 32-2 receives the 32-bit input data buses 36-2 and 37-2 and the iteration control register (ICR) line 33-2. Processor 32-2 provides 32-bit output on the bus 35-2 which connects to the GPR bus 65 and to the row 237-2 (dmc 2) of multiconnect 34.

The processor 32-3 includes a data memory1 (Mem1) functional unit 129 which receives input data on 32-bit bus 36-3 for storage at a location specified by the address bus 47-1. Processor 32-3 also provides output data at a location specified by address bus 47-1 on the 32-bit output bus 35-3. The output bus 35-3 connects to the GPR bus 65 and the multiconnect row 237-3 (dmc 3) of multiconnect 34. The Mem1 unit 129 connects to port (1) 153-1 for transfers to and from main store 7 and unit 129 has the same program (address space as distinguished from multiconnect addresses) as the main store 7. The processor 32-2 also includes a control (STUFF) functional unit 128 which provides an output on bus 35-5 which connects as an input to the predicate ICR multiconnect 29.

The processor 32-4 is the data memory2 (Mem2). Processor 32-4 receives input data on 32-bit bus 36-4 for

-27-

storage at an address specified by the address bus 47-2. Processor 32-4 also receives an ICR predicate input on line 33-4. Processor 32-4 provides an output on the 32-bit data bus 35-4 which connects to the GPR bus 65 and as an input to row 237-4 (dmc 4) of the multiconnect 34. Processor 32-4 connects to port(2) 153-2 for transfers to and from main store 7 and unit 32-4 has the same program address space (as distinguished from multiconnect addresses) as the main store 7.

The processing elements 32-1 through 32-4 have the input buses 36-1 through 36-4, 37-1 and 37-2 connected to a column of multiconnect elements 34, one from each of the row of elements 237-1 through 237-4 (dmc 1,2,3,4) as well as to an multiconnect element 34 in the GPR row 28 (mc0). Together, the rows 237-1 through 237-4 and a portion of the rows 28 and 29 form the data cluster multiconnect array 30.

The ICR multiconnect elements 22, including ICR(1), ICR(2), ICR(3), ICR(4) and the GPR(1), GPR(2), GPR(3), GPR(4), GPR(5), and GPR(6) multiconnect elements 34 are within the data cluster multiconnect array 30.

In FIG. 3, the processing elements 32-5, 32-6, and 32-9 form the address cluster of processing elements. The processor 32-9 is a displacement adder which adds an address on address bus 36-5 to a literal address on bus 44 from the I unit 32-7 to form an output address on bus 47-1 (amc6).

The processor 32-5 is address adder1(AAd1). The processor 32-5 receives an input address on bus 36-6 and a second input address on bus 37-3 and an ICR value from line 33-7. The processor 32-5 provides a result on output bus 47-3 which connects to the GPR bus 65 and to the row 48-2 (amc5) of multiconnect elements 34.

The processor 32-6 includes an address adder2(AAd2) functional unit and a multiplier(AMpy) functional unit which receive the input addresses on buses 36-7 and 37-4 and the ICR input on line 33-8. Processing element 32-6

-28-

provides an output on bus 47-4 which connects to the GPR bus 65 and to the row 48-1 (amc5) of the multiconnect elements 34.

In FIG. 3, the address adder1 of processor 32-5, performs three operations, namely, add (AAAd1), subtract (ASub1), and noop. All operations are performed on thirty two bit two's complement integers. All operations are performed in one clock. The operation specified is performed regardless of the state of the enable bit (WEN line 96, FIG. 13); the writing of the result is controlled by the enable bit.

The Address Adder 32-5 adds (AAAd1) the two input operands and places the result on the designated output bus 47-3 to be written into the specified address multiconnect register of row 48-2 (amc5) or into the specified General Purpose register of row 28 (mc0). Since the add operation is commutative, no special ordering of the operands is required for this operation.

The address subtract operation is the same as address add, except that one operand is subtracted from the other. The operation performed is operand B - operand A.

In FIG. 3, the Address Adder2 (AAAd2) 32-6 is identical to Address Adder1 32-5 except that adder 32-6 receives a separate set of commands from the instruction unit 32-7 and places its result on row 48-1 (amc5) of the Address Multi-Connect array 31 versus row 48-2, amc6.

In FIG. 3, the address adder/multiplier 32-6 performs three operations, namely, add (AAAd2), multiply (AMpy), and noop. All operations are performed on thirty two bit two's complement integers. All operations are performed regardless of the state of the enable bit; the writing of the result is controlled by the enable bit.

The Address Multiplier in processor 32-6 will multiply the two input operands and place the results on the designated output bus to be written into the specified row 48-1 (amc6) address multiconnect array 31 or into the specified

-29-

General Purpose Register row 28 (mc0). The input operands are considered to be thirty-two bit two's complement integers, and an intermediate sixty-four bit two's complement result is produced. This intermediate result is truncated to 31 bits and the sign bit of the intermediate result is copied to the sign bit location of the thirty two bit result word.

Each of the processing elements 32-1 through 32-4 in FIG. 3 is capable of performing one of the operations O_n^l where l designates the particular instruction in which the operation to be performed is found. The n designates the particular one of the operations. For example, the floating point add (FAdd) in processor 32-1 is operation $n=1$, the arithmetic and logic operation is operation $n=2$, and so on. Each operation in an instruction l commences with the same clock cycle. However, each of the processors for processing the operations may require a different number of clock cycles to complete the operation. The number of cycles required for an operation is referred to as the latency of the processor performing the operation.

In FIG. 3, the address multiconnect array 31 includes the rows 48-1 (amc6) and 48-2 (amc5) and a portion of the multiconnect elements in the GPR multiconnect 28 and the ICR multiconnect 29.

In FIG. 3, the instruction unit 32-7 has an output bus 54 which connects with different lines to each of the other processors 32 in FIG. 3 for controlling the processors in the execution of instructions. The processor 32-7 also provides an output on bus 35-6 to the GPR bus 65. Processor 32-7 connects to port(0) 153-0 for instruction transfers from main storage.

In FIG. 3, the processor 32-8 is a miscellaneous register file which receives the input lines 33-5 from the GPR(7) multiconnect element 34 and the line 33-6 from the ICR(5) element 22-5. The processor 32-8 provides an output on bus 35-7 which connects to the GPR bus 65.

-30-

The multiconnect arrays 30 and 31 consist of rectangular arrays of identical memory elements 34. Each multiconnect 34 is effectively a 64-word register file, with 32 bits per word, and is capable of writing one word and reading one word per clock cycle. Each row receives the output of one of the processors 32 and each column provides one of the inputs to the processors 32. All of the multiconnect elements 34 in any one row store identical data. A row in the multiconnect arrays 30 and 31 is effectively a single multi-port memory element that, on each cycle, can support one write and as many reads as there are columns, with the ability for all the accesses to be to independent locations in the arrays 30 and 31.

Each multiconnect element 34 of FIG. 3 contains 64 locations, each 32 bits wide. Specifying an address for a multiconnect element consists of specifying a displacement (via an offset field in the instruction word) from the location pointed to by a multiconnect pointer register (mcp) contained in each element 34 (register 82 of FIGS. 3 and 12). This mcp register can be decremented by 1 by the branch-to-top-of-loop operation controlled by a "Brtop" instruction.

Each element 34 in one preferred embodiment is implemented in two physical multiconnect gatearrays (67 and 68 in FIG. 11). Each gatearray contains 64 locations, each 17 bits wide corresponding to two bytes of the 4 byte multiconnect word (32 bits plus parity). Two read and one write addresses are provided in each cycle. Each physical gatearray supplies one half of the word for two logical elements 34.

The write address in register 75 of FIG. 12 for each element 34 is the location that will be written into each element for that row. All the multiconnect elements 34 in that row will receive the same write address. The Write address is stored in each element 34 and is not altered before it is used to write the random access memory in the

-31-

gatearray. The write data is also stored in register 73 and is unaltered before writing into the gatearray RAM 45 and 46.

The Read addresses are added in array adders (76 and 77 of FIG. 12 to the present value of the mcp. Each multiconnect element 34 contains a copy of the mcp in register 82 that can be either decremented or cleared. The outputs of the array adders are used as addresses to the gatearray RAM 45 and 46. The output of the RAM is then stored in registers 78 and 79.

Each multiconnect element 34 completes two reads and one write in one cycle. The write address and write data are registered at the beginning of the cycle and the write is done in the first part of the cycle and an address mux 74 first selects the write address from register 75.

After the write has been done, the address mux 74 is switched to the Aadd from adder 76. The address for the first or "A" read is added to the current value of the mcp to form Aadd(0:5).

The address selected by mux 74 from adder 77 for the second or "B" read is added to the current value of the mcp to form Badd(0:5). The A read data is then staged in a latch 89. Then the B read data and the latched A read data are both loaded into flip-flops of registers 78 and 79.

The address cluster 26 operates only on thirty-two bit two's complement integers. The address cluster arithmetic units 32-9, 32-5 and 32-6 treat the address space as a "circular" space. That is, all arithmetic results will wrap around in case of overflow or precision loss. No arithmetic exceptions are generated. The memory ports will generate an exception for addresses less than zero.

The address multiconnect array 31 of FIG. 3 is identical to the data multiconnect array 30 of FIG. 3 except for the number of rows and columns of multiconnect elements 34. The address multiconnect array 31 contains two rows 48-1 and 48-2 and six columns. Conceptually, each element

-32-

consists of a sixty-four word register file that is capable of writing one word and reading one word per clock. In any one row, the data contents of the elements 34 are identical.

All multiconnect addressing is done relative to the multiconnect pointer (mcp), except for references to the General Purpose Register file 28. The multiconnect pointer (mcp) is duplicated in each multiconnect element 34 in a mcp register 82 (see FIG. 12). This 6-bit number in register 82 is added in adders 76 and 77 to each register address modulo 64. In the example described, The mcp has the capability of being modified (decremented) and of being synchronized among all of the copies in all elements 34. The mcp register 82 (see FIGS. 6 and 12) is cleared in each element 34 for synchronization. However, for alternative embodiments, synchronization of the mcp registers is not required.

The General Purpose Register file is implemented using the multiconnect 28 row of elements 34 (mc0). The mcp for the GPR 28 is never changed. Thus, the GPR is always referenced with absolute addresses.

The value of the mcp at the time the instruction is issued by instruction unit 32-7 is used for both source and destination addressing. Since the destination value will not be available for some number of clocks after the instruction is issued, then the destination physical address must be computed at instruction issue time, not result write time. Since the source operands are fetched on the instruction issue clock, the source physical addresses may be computed "on the fly". Since the mcp will be distributed among the multiconnect elements 34, then each multiconnect element provides the capability of precomputing the destination address, which will then be staged by the various functional units.

The destination address is added to mcp only if the GIB select bit is true. The GIB select bit is the most

-33-

significant bit, DI(6) on line 64-1 of FIG. 4, of the seven bit destination address DI on bus 64. If the GIB select bit is false, then the destination address is not added to mcp, but passes unaltered.

Certain operations have one source address and two for destination addresses. For these instances, the value of mcp is connected from the multiconnect element 34 via line 50 of FIG. 12 so that its value may be used in external computations. Bringing mcp off the chip also provides a basis for implementing logic to ensure that the multiple copies of mcp remain synchronized.

Instruction Unit-FIG. 4

Further details of the instruction unit (IU) of FIG. 3 are shown. The instruction unit includes an instruction sequencer 51 which provides instruction addresses, by operation of address generator 55, to the instruction memory 52. Instruction memory 52 provides an instruction into the instruction register 53 under control of the sequencer 51. For a horizontal architecture, the instruction register 53 is typically 256 bits wide. Register 53 has outputs 54-1 through 54-8 which connect to each of the processing elements 32-1 through 32-8 in FIG. 3.

Each of the outputs 54 has similar fields and includes an opcode (OP) field, a first address source field (S1), a second address source field (S2), a destination field (D1), a predicate field (PD), and a literal field (LIT). By way of example, the output 54-2 is a 39-bit field which connects to the processor 32-2 in FIG. 3. The field sizes for the output 43-2 are shown in FIG. 4.

While the field definition for each of the other outputs 54 from the instruction register 53 are not shown, their content and operation is essentially the same as for output 54-2. The instruction unit bus 54-8 additionally includes a literal field (LIT) which connects as an input to bus 44 to the displacement adder 32-9 of FIG. 4.

-34-

The instruction unit 33-7 of FIGS. 3 and 4 provides the control for the computer 3 of FIG. 3 and is responsible for the fetching, caching, dispatching, and reformatting of instructions.

The instruction unit includes standard components for processing instructions for controlling the operation of computer 3 as a pipelined parallel processing unit. The following describes the pipeline structure and the operation of the address generator 55.

The pipeline in address generator 55 includes the following stages:

| C | I | E1 ... En | D |

- C: During the C Cycle the ICache 52 of FIG. 4 is accessed. At the end of this cycle, the instruction register 53 is loaded.
- I: During the I Cycle, the Instruction Register 53 is valid. The Opcodes, Source (S1, S2), Destination (D1) and other fields on buses 54 are sent to the various processors 32. The Source fields (S1, S2) access the multiconnects 34 during this cycle.
- E: The E cycle or cycles represent the time that processors 32 are executing. This E cycle period may be from 1 to n cycles depending on the operation. The latency of a particular operation for a particular processor 32 is $(n + 1)$, where n is the number of E cycles for that processor.
- D: During the D cycle, the results of an operation are known and are written into a target destination. An instruction that is in an I cycle may access the

- 35 -

results that a previous instruction provided in its D cycle.

In multiple-operation (MultiOp) mode, which occurs when register 53 of FIG.4 is employed, there are up to seven operations executing in parallel. This parallel operation is shown as:

C	I	E1	E2	E3	D				
		E1	E2	E3	E4	E5	E6	D	
		E1	E2	E3	E4	E5	D		
		E1	E2	E3	D				
		E1	E2	E3	E4	D			
		E1	E2	E3	D				
		E1	E2	E3	E4	D			

The following is an example of the operation of a sequence of instructions starting at address A. the CurrIA address is used to access the ICache 52 of FIG. 4.

	0	1	2	3	4	5	6	7	8	
A: I0	T	C	I	E1	D					Latency = 2
A+1: I1		T	C	I	E1	E2	D			Latency = 3
A+2: I2			T	C	I	E1	E2	E3	D	Latency = 4
A+3: I3				T	C	I	E1	E2	D	Latency = 3

CurrIA		A	A+1	A+2	A+3			
IR53			I0	I1	I2	I3		

The following is an example of the operation of a Branch Instruction. Whether the branch instruction is conditional or unconditional does not matter. During the first half of Cycle 4, the Branch address is calculated. During that cycle the Tag and the TLB Arrays are accessed by the sequential address (A+3) in the first half, and by the

-36-

Branch Address (T) in the second half. Also during Cycle 4, a Branch Predicate is accessed.

In Cycle 5 both the location in the ICache of the Sequential Instruction and the Target Instruction are known. Also known is the branch condition. The branch condition is used to select between the Sequential Instruction address and the Target Instruction address when accessing the ICache 52. If the Branch is an unconditional Branch, then the Target Instruction will always be selected.

	1	2	3	4	5	6	7	8	9
A: I0	T	C	I						
A+1: Br T		T	C	I					
A+2: I2			T	C	I				
A+3: I3				T					
T: J5				T	C	I			
T+1: J6					T	C	I		
NextIA	A	A+1	A+2	A+3	A+4	T+2			
CurrIA		A	A+1	A+2	A+3	T+1			
NextTIA					T+1				
CurrTIA					T				
IR			I0	BrT	I2	J5	J6		

The Timing and Loop Control 56 of FIG. 4 is control logic which controls the Iteration Control Register (ICR) multiconnect 29 in FIG. 3 and FIG. 14, in response to Loop Counter 90, Multiconnect/ICR Pointer Registers (mcp 82 in FIG. 12 and icp 102 in FIG. 14), Epilog Stage Counter 91.

-37-

Control 56 is used to control the conditional execution of the processors 32 of FIG. 3.

The control 56 includes logic to decode the "Brtop" opcode and enable the Brtop executions. The control 56 operates in response to a "Brtop" instruction to cause instruction fetching to be conditionally transferred to the branch target address by asserting the BR signal on line 152. The target address is formed by address generator 55 using the sign extended value in the "palit" field which has a value returned to the sequencer 51 on bus 54-8 from the instruction register 53 and connected as an input on line 151 to address generator 55 in FIG. 4 and FIG. 5.

The Loop Counter (lc) 90, Epilog Stage Counter (esc) 91, and the ICR/Multiconnect Pointers (icp/mcp) register 82 of FIG. 6, FIG. 12 and 102 of FIG. 14 are conditionally decremented by assertion of MCPEN and ICPEN from control 56 in response to the Brtop instruction. The "icr" location addressed by $(icp - 1) \bmod 128$ is conditionally loaded in response to the Brtop instruction into register 92 of FIG. 14 with a new value in response to the signals on lines 104 from control 56.

The branch latency and the latency of the new values in the "lc", "esc", "icr", and "icp/mcp" are 3 cycles after "Brtop" is issued. Two additional instructions execute before the "Brtop" operation takes effect.

The "lc" value in register 90 and "esc" in register 91 of FIG. 4 are checked by control 56 on the cycle before the "Brtop" is to take effect (latency of 2 cycles) to determine what conditional operations should occur.

The control 56 operates in the following manner in response to "Brtop" by examining "lc" on 32-bit bus 97 and "esc" on 7-bit bus 98. If the "lc" is negative, the "esc" is negative, or if the "lc" and "esc" are both zero, then the branch is not taken (BR not asserted on line 151); otherwise, the branch is taken (BR asserted on line 151).

-38-

If the "lc" is greater than zero, then the "lc" is decremented by a signal on line 257; otherwise, it is unchanged.

If the "lc" is zero, and the "esc" is greater than or equal to zero, then the "esc" is decremented by a signal on line 262; otherwise, it is unchanged.

If the "lc" is positive, and the "esc" is greater than or equal to zero, then the "icp/mcp" is decremented by generating MCPEN and ICPEN on lines 85 and 86.

The Iteration Control Register(icr) 92 of FIG. 14 is used to control the conditional execution of operations in the computer 3. Each "icr" element 22 in FIG. 2 and 92 in FIG. 14 consists of a 128 element array with 1 bit in each element. On each cycle, each "icr" element 22 can be read by a corresponding one of the seven different processors (FMpy) 30-2, (FAdd) 30-1, (Mem1) 32-3, (Mem2) 32-4, (AAd1) 32-5, (AAd2) 32-6, and (Misc) 32-8. Each addressed location in the "icr" 92 is written implicitly at an "icr" address in response to the "Brtop" instruction.

An "icr" address is calculated by the addition of the "icrpred" field (the PD field on the 7-bit bus 71 of FIG. 4, for example) specified in an NP operation with the "ICR Pointer" (icp) register 102 at the time that the operation is initiated. The addition occurs in adder 103 of FIG. 14.

The Loop Counter "lc" 90 in FIG. 4 is a 32-bit counter that is conditionally decremented by a signal on line 257 during the execution of the "Brtop" instruction. The loop counter 90 is used to control the exit from a loop, and determine the updating of the "icr" register 92.

The Epilog Stage Counter "esc" 91 is a 7-bit counter that is conditionally decremented by a signal on line 262 during the execution of the "Brtop" instruction. The Epilog Stage Counter 91 is used to control the counting of epilog stages and to exit from a loop.

The detailed logical statement of the logic for control 56 for controlling the "lc" counter, the "esc"

-39-

counter in response to "Brtop" appears in the following CHART.

CHART

```

if ![(lc@2<0) :: (esc@2<0) :: (lc@2==0) && (esc@2==0)]
    pc@3 = brb@0 + paLit;
if (lc@2>0) lc@3 = lc@2 - 1;
if [(lc@2==0) && (esc@2>=0)] esc@3 = esc@2 - 1;
if [(lc@2>=0) && (esc@2>=0)]
    {icr[icp@2 - 1]@3 = [(lc@2>0)? 1 : 0];
    icp@3 = icp@2 - 1; mcp@3 = mcp@2 - 1;}

```

!	means	NOT
::	means	OR
&&	means	AND
==	means	COMPARE FOR EQUAL TO
>=	means	COMPARE FOR GREATER OR
		EQUAL TO
=	means	SET EQUAL TO
@	means	OFFSET TIME
>	means	GREATER THAN

Instruction Address Generator-FIG. 5

In FIG. 5, further details of the instruction address generator 55 of FIG. 4 are shown. The generator 55 receives an input from the general purpose register file, GPR (7) via the input bus 33-5. The bus 33-5 provides data which is latched into a branch base register (BRB) 205. Register 205 is loaded as part of an initialization so as to provide a branch base address. The BRB register 205 provides an input to a first register stage 144-1 which in turn connects directly to a second register stage 144-2. The output from the register 144-2 connects as one input to the adder 146.

In FIG. 5, the address generator 55 receives the literal input (palit) on bus 151 which is derived through the

-40-

timing loop control 56 of FIG. 4 directly from the instruction register 53 via the bus 54-8. In FIG. 5, the bus 151 has the literal field latched into the first register stage 145-1, which in turn is connected to the input of the second register stage 145-2. The output from the second register stage 145-2 connects as the second input to adder 146. The adder 146 functions to add a value from the general purpose register file, GPR (7), with a literal field from the current instruction to form an address on bus 154. That address on bus 154 is one input to the multiplexer 148. Multiplexor 148 receives its other input on bus 155 from the address incrementer 147. Incrementer 147 increments the last address from the instruction address register 149. The multiplexer selects either the branch address as it appears on bus 154 from the branch adder 146 or the incremented address on the bus 155 for storing into the instruction address register 149. The branch control line 152 is connected as an input to the multiplexer 148 and, when line 152 is asserted, the branch address on bus 154 is selected, and when not asserted, the incremented address on bus 155 is selected. The instruction address from register 149 connects on bus 150 as an input to the instruction cache 52 of FIG. 4.

The registers 144-1, 144-2 together with the instruction address register 149, and the registers 145-1, 145-2 together with the instruction address register 149, provide a three cycle latency for the instruction address generator 55. In the embodiment described, the earliest that a new branch address can be selected for output on the bus 150 is three cycles delayed after the current instruction in the instruction register 53 of FIG. 4. Of course, the latency is arbitrary and may be selected at many different values in accordance with the design of the particular pipeline data processing system.

Invariant Addressing Unit-FIG. 6

-41-

The invariant addressing unit 12 in FIG. 6 is typical of each of the units 12 in FIG. 2 and each includes a modifying unit 76, such as subtracter 84, for forming a current pointer address, $mcp(\bar{i})$, from a previous pointer address $mcp(\bar{i}-1)$ with the operation $D^*[mcp(\bar{i}-1)]$ such that $mcp(\bar{i}) = D^*[mcp(\bar{i}-1)]$. The unit 12 includes a pointer register 82 for storing the pointer address, $mcp(\bar{i})$, for use in the \bar{i}^{th} -iteration. The unit 12 includes an address generator (adder 76) combining the pointer address, $mcp(\bar{i})$, with an address offset, $ao_n^k(c)$, to form the memory addresses, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration, which address is connected to memories 34 to provide an output on the c^{th} port. For the invariant address units 12-1 in FIG. 2, the c^{th} port ($c=1$) is ports 14-1 and 10-1. For the Address units 12-2 the c^{th} port ($c=2$) is ports 14-2 and 10-2. Processor 32-1 has first and second input ports ($c=1, c=2$) and the other processors have one or more similar inputs ($c=3, c=4, \dots$).

Typical Processor-FIG. 7

In FIG. 7, further details of a typical one of the processors 32 of the FIG. 2 system are shown. The processor 32 includes one or more functional units 130. In FIG. 7, the functional units include the functional units 130-1 and 130-2. The functional units include well-known execution devices, such as adders, multipliers, dividers, square root units, arithmetic and logic units, and so forth. Additionally, in accordance with the present invention, the functional units also include data memories for storing data. When the functional units 130-1 and 130-2 of FIG. 7 perform arithmetic and logical functions, the functional units typically include first and second inputs, namely input bus 36 and input bus 37. The buses 36 and 37 are the data buses which carry data output from the multiconnect array of FIG. 2. Each functional unit 130 includes a number of shift-register stages (first-in/first-out stack), x , which represents the latency time of the

-42-

functional unit, that is, the number of cycles required for the input data on buses 36 and 37 to provide valid outputs on buses 35, including the bus 35-1 from the unit 130-1 and the bus 35-2 from the unit 130-2. The number of stages 132-1, 132-2, ..., 132-x determining the latency time is a variable and the different processors 32 of FIG. 2 may each have different number of stages and latency times. Similarly, each functional unit 130-1 and 130-2 within a processor may have different latency times. For example, the functional unit 130-1 has a latency of x and the functional unit 130-2 has a latency of y. The functional unit 130-2 has the stages 132-1, 132-2, ..., 132-y which operate as a first-in/first-out stack.

In FIG. 7, an opcode decoder 137 receives the opcode on bus 63 from the instruction register 53 of FIG. 4. Decoder 137 provides a first output on line 156 for enabling the functional unit 130-1 and provides a second output on line 157 for enabling the second functional unit 130-2. Similarly, the enable signals from decoder 137 are input on lines 156 and 157 to the processor control 131.

In FIG. 7, a predicate stack 140 receives the predicate line 33 from one of the ICR registers 22 of FIG. 3. The predicate stack 140 includes a number of stages 140-1, 140-2, ..., 140-x,y which is equal to the largest of x or y. When functional unit 130-1 is employed, the predicate stack utilizes x stages and, when functional unit 130-2 is enabled, the predicate stack 140 employs y stages so that the latency of the predicate stack matches that of the active functional unit.

Each of the stages in the stack 140 provides an input to the control 131. In this manner, the control 131 is able to control the operation and the output as a function of the predicate bit value in any selected stage of the stack 140.

In FIG. 7, the processor 32 includes an address first-in/first-out stack 133. The address stack receives

-43-

the $D1_{in}$ bus 164 from the instruction register 53 of FIG. 4. The instruction stack 133 includes the largest of x or y stages, namely 133-1, 133-2, ..., 133- x,y . Whenever the functional unit 130-1 is enabled, x stages of the stack 133 are employed and the output 264 has latency x under control of line 265 from control 131, and whenever the functional unit 130-2 is enabled, y stages of the stack 133 are employed and the output 264 has latency y under control of line 265. The processing unit 32 of FIG. 7 operates such that the latency of the particular unit enabled and the latency of the predicate stack 140 and the address stack 133 are all the same. In this manner, the pipeline operation of the processing units are maintained synchronized. In the simplest example, the processor 32 only need include a single functional unit having a single latency x for the functional unit 130-1, the predicate stack 140 and the address stack 133. The inclusion of more than one functional unit into a single processor is done for cost reduction.

In FIG. 7, the control unit 131 receives inputs from the decoder 137, the predicate stack 140 and the functional units 130-1 and 130-2. The control unit 131 provides a write enable (WEN) signal on line 96. The write enabled signal on line 96 can be asserted or not asserted as a function of the state of a predicate bit and/or as a function of some condition created in a functional unit 130-1 or 130-2. The write enable signal on line 96 connects to the multiconnect 30 of FIG. 2 and determines when the result on bus 35-1 or 35-2 is actually to be written into the respective row of multiconnect elements.

STUFFICR Processor-FIG. 8

In FIG. 8, further details of the STUFF processor 32-3, one of the processors 32 of FIG. 2, are shown. In FIG. 8, the processor 32-3 includes a functional unit 130-3, which has a latency of three cycles. The three

-44-

cycles are represented by the register stages 158-1, 158-2, and 158-3. The input to register stage 158-1 is from the column data bus 36-3. A comparator 159 receives the output from stage 158-2 and compares it with a "0" input. If the input operand on bus 36-3 is all 0's, then the output from comparator 159 is asserted as a logical 1 connected to the EXCLUSIVE-OR gate 160. The other input to gate 160 is derived from the opcode decoder 137-3. The opcode decoder 137-3 functions to decode the opcode to detect the presence of either a STUFFICR or a STUFFBAR (STUFFICR̄) opcode. Whenever the STUFFICR opcode is decoded by decoder 137-3, the signal on line 168 is asserted and latched into stage 161-1 and the next clock cycle is latched into stage 161-2 to provide an input to the EXCLUSIVE-OR gate 160. During the same clock cycles, the predicate bit from line 33-3 is latched into the stage 163-1 if AND gate 180 is enabled by a decode (indicating either STUFFICR or STUFFBAR) of decoder 137-3 on line 181. In the next cycle the data in stage 162-1 is latched into the stage 163-2 to provide an input to the AND gate 162. The output EXCLUSIVE-OR gate 160 forms the other input to AND gate 162. The output from gate 162 is latched into the register 158-3 to provide the ICR data on line 35-5 written into the addressed location of all predicate elements 22 in predicate multiconnect 29.

In FIG. 8, the predicate address on bus 164-3 derives from the predicate field (PD) as part of bus 54-8 from the instruction register 53 of FIG. 4, through the invariant address unit 12-3 in FIG. 2 and is input to the address stack 133-3. The predicate address is staged through the stages 165-1, 165-2 and 165-3 to appear on the predicate address output 264-3. The predicate address on bus 264-3 together with the WEN signal on line 96-3 addresses the row of ICR multiconnect elements 22 to enable the predicate bit on line 35-5 to be stored into the addressed location in each element 22.

-45-

In FIG. 8, the latency of the functional unit 130-3, the control 131-3, the predicate stack 140-3, and the address stack 133-3 are all the same and equal three cycles.

Iselect-Multiply Processor-FIG. 9

In FIG. 9, further details of the processor 32-2 of FIG. 3 are shown. The functional unit 130-4 includes a conventional multiplier 169 which is used to do multiplies, divides and square roots as specified by lines 267 from decoder 137-2. Additionally, either the data input on bus 36-2 or the data input on bus 37-2 can be selected for output on the bus 35-2. The selection is under control of a predicate bit from the predicate stack 140-2 on line 266 to multiplexer 171.

In FIG. 9, the bus 36-2 connects through the left register stage 170-1 as one input to the multiplier 169 including stages 170-2, ..., 170x. The bus 37-2 connects through the right register stage 170-1 as the other input to the multiplier 169. Additionally, the outputs from the registers in stage 170-1 connect as inputs to the multiplexer 171. Multiplexor 171 selects either the left or right register from the stage 170-1 to provide the data onto the bus 276 as one input to the multiplexer 172 through register stack 170. Multiplexor 171 is controlled to select either the left or right input, that is the latched value of the data from bus 36-2 or from bus 37-2, under control of a predicate bit latched in the stage 174-1. The predicate latched into stage 174-1 is a 1 or 0 received through the AND gate 268 from the predicate line 33-2 which connects from the ICR multiconnect element 22-2 of FIG. 3. The AND gate 268 is enabled by a signal on line 269 asserted by decoder 137-2 when an Isel operation is to occur and multiplier 169 is to be bypassed. Also, gate 268 is satisfied when a multiply or other function is to occur with multiplier 169 and the value of the predicate bit on line 33-2 will be used, after propagation to stage 174-x1

-46-

to control the enable or disable of the storage of the results of that operation.

In FIG. 9, the multiplier 169 combines the input operands from the register stages 171-1 and processes them through a series of stages 170-2 through 170-x. The number of stages x ranges from 1 to 30 or more and represents the number of cycles required to do complicated multiplications, divisions or square root operations. The same number of stages 170'-2 to 170'-x connect from multiplexer 171 to multiplexer 172. The output selected from the 170-x and 170'-x stages connects through the multiplexer 172 to the final stage 170-x1. Multiplexor 172 operates to bypass the multiplier functional unit 169 whenever an iselect, Isel, opcode is detected by decoder 137-2. The decoder 137-2 decodes the opcode and asserts a signal which is latched into the register stage 176-1 and transferred through stages 176-2 to 176-x. When latched in stage 176-x, the multiplexer 172 is conditioned by the signal on line 265 to select 170'-x as the input to register 170-x1. Otherwise, when a multiply or other command using the multiplier 169 is decoded, multiplexer 172 selects the output from stage 170-x for latching into the stage 170-x1. Whenever an Isel command is asserted, the register 176-x1 stores a 1 at the same time that the selected operand is output on the bus 35-2. The 1 in register 176-x1 satisfies the OR gate 177, which in turn enables the write enable signal, WEN, on line 96-2. The WEN signal on line 96-2, together with the destination address on bus 264-2, is propagated to multiconnect 237-2 (dmc2) to store the data on bus 35-2 in each element 34 of the row.

When the decoder 137-2 does not detect an iselect command, then the OR gate 177 is satisfied or not as a function of the 1 or 0 output from the predicate stack stage 174-x1. The number of stages in the predicate stack 140-2 includes 174-1, 174-2, ..., 174-x, and 174-x1. Therefore, the latency of the predicate stack 140-2 is the same as the

-47-

latency of the functional unit 130-4 when the multiply unit 169 is employed. When the iselect command is present, then the latency for the write enable signal WEN is determined by the delays 176-1 to 176-x1 which matches the latency of the operand through the multiplexer 171 and 172 bypass in the functional unit 130-2. The input stage 176-1 of stack 176 is loaded with a 1 whenever an Isel operation is decoded and is otherwise loaded with a 0. Therefore, OR gate 177 will always be satisfied by a 1 from stage 176-x for Isel operations. However, for other operations, the predicate value in stage 174-x1 will determine whether gate 177 is satisfied to generate the WEN signal.

In a similar manner, the address stack 133-2 has a latency which is the same as the functional unit 130-2, both under the condition of the latency through the multiplier 169 or the latency through the multiplexers 171 and 172. As an alternative, stacks 170', 176, and 178' may have a different number of stages and latency than stacks 170, 174 and 178.

In FIG. 9, the multiplexer 179, like the multiplexer 172, bypasses the register stages 178-2 through 178-x when the Isel opcode is decoded by decoder 137-2. When the Isel command is not decoded, then multiplexer 179 selects the output from stage 178-x as the input to stage 178-x1. In this way, the latency of the address stack 133-2 remains the same as for the functional unit 130-4.

The Isel structure of FIG. 9 (MUX's 171, 172,...) is typical of the select structure in each processor 32 of FIGS. 2 and 3. The select employed in floating point processor 32-1 is identified as "Fsel", for example.

Processing Element Multiconnect-FIG. 10

In FIG. 10, the multiconnect elements 34 are organized in an array 30 corresponding to a portion of the data cluster of FIG 3. Only the detail is shown for the processor 32-2 of FIGS. 2, 3 and 9 and the corresponding

-48-

multiconnect elements 34. In particular, the third column of multiconnect elements 34 designated D(1,3), D(2,3), D(3,3), and D(4,3), all have an output which connects in common to the 33-bit first data in (DI1) bus 36-2. Additionally, the general purpose register element (GPR)(3) has an output which connects to the bus 36-2. Each of these elements is similarly addressed by the 9-bit S1 source bus 59 through the invariant address units 12.

In FIG. 10, the fourth column of multiconnect elements 34 includes the elements GPR(4), D(1,4), D(2,4), D(3,4), and D(4,4). The output of each of these elements in column four connects to the 33-bit second data input (DI2) bus 37-2. Also each of the multiconnect elements in column four is addressed by the 9-bit second source S2 bus 61 through an invariant address unit 12. Data from anyone of the column three multiconnect elements addressed by the S1 bus 59 provides data on the DI1 bus 36-2. Similarly data from any one of the multiconnect elements in column four is addressed by the S2 bus to provide data on the DI2 bus 37-2. The processor (PE) 32-2 performs an operation on the input data from buses 36-2 and 37-2 under control of the opcode on bus 63. Bus 63 is the OP(6:0) field from the output 54-2 of the instruction register 53 of FIG 4. The operation performed by the processor 32-2 has a result which appears on the data outbus 35-2. The data outbus 35-2 connects as an input to each of the multiconnect elements comprising the dmc2 row of the data cluster. The dmc2 row includes the multiconnect elements D(3,1), D(3,2), D(3,3), D(3,4), and D(3,5).

In FIG. 10, the destination address appears on the D1 bus 64 which is derived from one of the fields in the output 54-2 from the instruction register 53 of FIG 4. The bus 64 connects to invariant address unit 12'-12 forming the address on bus 164. The address is delayed in address stack 1 to provide the destination address on bus 264. The data output bus 35-2 also connects in common to the GPR bus

-49-

65 which forms a data input to the mc0 row of GPR elements, GPR(1) to GPR(13) which form the GPR multiconnect 49.

In FIG. 10, the destination address output from processing element 32-1 on line 264-1, together with the write enable, WEN, line 96-1, form the bus 271 which are connected to each of the elements in the dmc1 multiconnect comprised of the elements D(4,1), ..., D(4,5). Also, the line 96-1 and the bus 264-1 connect to the common GPR destination address bus 270 which connects to address, when enabled, each of the elements in the GPR multiconnect 49. In a similar manner, the processing element 32-2 has its write enable, WEN, line 96-2 connected together with the destination address bus 264-2 to form the bus 273 which addresses the dmc2 row of multiconnect elements 34, including the elements D(3,1), ..., D(3,5). The line 96-2 and the bus 264-2 also connect to the common bus 270 which provides the destination address and enable to the GPR multiconnect 49. The destination address bus 270 connects in common to all of the processing elements, such as processing elements 32-1 and 32-2, but only one of the destination addresses is active at any one time to provide an output which is to be connected in common to the GPR multiconnect 49. As indicated in FIG. 9, the WEN enable signal on line 96-2 enables the outgating of the registers 170-x1 and 178-x1 which provide the data output on bus 35-2 and the destination address output on bus 264-2. This gating-out of the registers ensures that only one element will be connected to the common bus 65 and one element to the common bus 270 for transmitting destination addresses and data to the GPR multiconnect 49. In a manner similar to FIG. 9, all of the other processing elements 32 of FIG. 2 and FIG. 3 which connect to the GPR bus 65 and the corresponding destination address bus 270 (see FIG. 10) are enabled by the respective write enable signal, WEN, to ensure that there is no contention for the common buses to the GPR multiconnect 49.

-50-

In FIG. 10, the pair of multiconnect elements D(3,3), and D(3,4) comprise one physical module 66. The combination of two pairs of logical modules, like modules D(3,3), and D(3,4) is arbitrary as any physical implementation of the multiconnect array maybe employed.

Physical Multiconnect Module FIG. 11

In FIG. 11, the module 66 is a typical implementation of two logical modules, such as D(3,3) and D(3,4) of FIG. 10. In FIG. 11, a first (C1) chip 67 and a second (C2) chip 68 together form the logical modules, D(3,3), and D(3,4), of FIG. 10. However, one half of the D(3,3) multiconnect element is in the C1 chip 67 and one half is in the C2 chip 68. Similarly, one half of the D(3,4) logical multiconnect appears in each C1 chip 67 and C2 chip 68. Both the chip 67 and 68 receive the S1 source bus 59 and the S2 source bus 61. The S1 source bus 59 causes chips 67 and 68 to each provide the 17-bit data output C1(AO) and C2(AO), respectively, on output lines 69 and 70. The outputs on lines 69 and 70 are combined to provide the 33-bit DI1 data bus 36-2.

In a similar manner, the address on the S2 address bus 61 addresses both the C1 chip 67 and the C2 chip 68 to provide the C1(BO), and the C2(BO) 17-bit outputs on lines 71 and 72, respectively. The data outputs on lines 71 and 72 are combined to form the 33-bit data DI2 on bus 37-2.

The DI1 and DI2 data buses 36-2 and 37-2 connect as the two inputs to the processor 32-2 and FIG. 10.

In FIG. 11, the D1 destination bus 273 connects as an input to both the C1 chip 67 and the C2 chip 68. The destination address and the WEN signal on the D1 bus 273 causes the data out data on the DO bus 35-2 to be stored into both the C1 chip 67 and the C2 chip 68.

Multiconnect Chip FIGS. 12 and 13

-51-

In FIGS. 12 and 13, the C1 chip 67 of FIG. 11 is shown as typical of the chips 67 and 68 and the other chips in each of the other multiconnect elements 34 of FIG. 10, taken in pairs.

In FIG. 12, two 64 x 10 random access memories (RAM'S) 45 and 46 are the data storage elements. The data out bus

-52-

line 21 is enabled only when the write enable (EN) signal on line 96, the signal on line 123 and the write strobe (WRSTRB) signal on line 124 are present. In the absence of any of these signals, the MEN signal on line 121 is not asserted and no write occurs into the RAMS 45 and 46.

In FIG. 13, the WEN signal on line 96 is generated by the corresponding processor, in the example being described, the processor 32-2, in FIG. 10. The processor 32-2 when it completes a task, provides an output on the output data bus 35-2 and generates the WEN signal unless inhibited by the predicate output on line 33-2. The WEN signal on line 96 is latched into the register 113 which has its inverted output connected to the ORgate 120.

In FIG. 13, the signal on line 123 is asserted provided that the row ID (ROWID) on line 125 is non-zero and provided the GPR register has not been selected as evidenced by the signal on line DI(6), line 64-1 is zero. Under these conditions, the line 123 is asserted and stored in the register 114. The double bar on a register indicates that it is clocked by the clock signal along with all the other registers having a double bar. If both registers 113 and 114 store a logical one, a logical zero on the strobe line 124 will force the output of gate 120 to a logical zero thereby asserting the MEN signal on line 121. If either of registers 113 or 114 stores a zero, then the output from gate 120 will remain a logical one and the MEN signal on line 121 will not be asserted.

In FIG. 13, the 3-bit input which comprises the ROWID signal on line 125 is hardwired to present a row ID. Each element 34 in FIGS. 3 and 10 is hardwired with a row ID depending on the row in which the element is in. All elements in the same row have the same ROWID.

In FIG. 13, comparator 118 compares the ROWID signal on line 125 with the high order 3-bit address S1(8:6) on line 59-1 from the instruction register 53 from FIG. 4.

-53-

If a comparison occurs, a one is stored into the register 116 to provide a zero and assert the A enable (AEN) signal on line 126. The AEN signal on line 126 connects to the ANDgate 80 to enable the output from the register 78 in FIG. 12.

In a similar manner, the comparator 119 compares the ROWID signal on line 125 with the three high order bits S2(8:6) on lines 61-1 from the instruction register 53 of FIG. 4. If a compare occurs, a one is clocked into the register 117 to enable the B enable signal (BEN) on line 127. The BEN signal on line 127 connects to the ANDgate 81 in FIG. 12 to enable the contents of register 79 to be gated out from the chip 67 of FIG. 11.

ICR ELEMENT FIG. 14

In FIG. 14, a typical one of the elements 22, namely element 22-2, which form the row 29 of ICR elements in FIG. 3 is shown. The ICR register 92 provides a 1-bit predicate output on line 33-3. The ICR register 92 is addressed by a 7-bit address from the adder 103. Adder 103 forms the predicate address by adding the offset address in the ICP pointer register 102 to the predicate (PD) address on the 7-bit bus 167 which comes from the instruction register 53 of FIG. 4 as connected through the processor 32-2 of FIG. 10.

The iteration control register (ICR) 92 can have any one of its 128 locations written into by the 1-bit ICR data (ICRD) line 108 which comes from the timing and loop control 56 of FIG. 4. The logical one or zero value on line 108 is written into the ICR 92 when the write interval control register (WICR) signal on line 107 is asserted. The enable signal on line 107 is derived from the timing and loop control 56 of FIG. 4. The address written into in register 92 is the one specified by the adder 103.

In FIG. 14, the ICP register 102 stores a pointer which is an offset address. The contents of register 102

-54-

are initially cleared whenever the (ICPCLR) line 109 from the timing and loop control 56 of FIG. 4 is asserted. When line 109 is asserted, the output of gate 106 is a zero so that when ICPEN is enabled the register 102 is clocked to the all zero condition. When the ICPCLR line 109 is not asserted, then the assertion of the enable signal ICPEN on line 110 causes register 102 to be incremented by one unit. In the embodiment described, the incrementing of the register 102 occurs by subtracting one in the subtracter 105 from the current value in register 102. The incrementing process is actually a decrementing of the contents of register 102 by 1.

Single And Multiple Operation Unit-FIG. 15

In FIG. 15, the single operation/multiple operation unit which forms a modification to the instruction unit of FIG. 4 is shown. In FIG. 4, the instruction register 53 is replaced by the entire FIG. 15 circuit. The input bus 184 from the instruction cache 52 of FIG. 4 connects, in FIG. 15, to the input register 53-1. Register 53-1 as an input receives information in the same way as register 53 of FIG. 4. The output from the unit of FIG. 15 from register 53-2 on the buses 54-1, 54-2, ..., 54-8 and is like the outputs from the register 53 in FIG. 4.

In FIG. 15, for multiple operations, the input register 53-1 is connected directly to the output register 53-2 through the multiplexers 190. Register 53-1 includes a stage for each operation to be executed, each stage including an opcode field, source and destination offset addresses, predicate fields and other fields as previously described in connection with FIG. 4. The output from each stage of register 53-1 appears on buses 193-1, 193-2, ..., 193-8, having the same information as buses 54-1, 54-2, ..., 54-8 of FIG. 4. These buses 193-1 through 193-8 in turn connect to the multiplexers 190-1, 190-2, ..., 190-8 which have outputs which connect in turn to the

-55-

corresponding stages of the output register 53-2 so as to directly provide the outputs 54-1, 54-2, ..., 54-8, respectively. The outputs from input register 53-1 are connected directly as inputs to the output register 53-2 when the control line 194 from the mode control register 185 is asserted to indicate a multiop mode of operation.

When the mode control 185 does not assert line 194, indicating a single operation mode, the multiplexers 190-1 through 190-8 are active to select outputs from the selector 188. Only one output from selector 188 is active at any one time, corresponding to the single operation to be performed, and the other outputs are all nonasserted.

Selector 188 derives the information for a single operation from the multiplexer 187. Selector 188, under control of the control lines 192, selects one of the multiplexers 190-1 through 190-8 to receive the single operation information from multiplexer 187. The particular one of the operations selected corresponds to one of the multiplexers 190-1 through 190-8 and a corresponding one of the output buses 54-1 through 54-8.

Multiplexor 187 functions to receive as inputs each of the buses 277-1 through 277-7 from the input register 53-1. Note that the number (7) of buses 277-1 through 277-7 differs from the 8 buses 190-1 to 190-8 since the field sizes for single operation instructions can be different than for multiple operation instructions. Multiplexor 187 selects one of the inputs as the output on buses 191 and 192. The particular one of the inputs selected by multiplexer 187 is under control of the operation counter 186. Operation counter 186 is reset each time the control line 194 is nonasserted to indicate loading of single operation mode instructions into register 53-1 and register 185. Thereafter, the operation counter 187 is clocked (by the system clock signal not shown) to count through each of the counts representing the operations in input register 53-1.

-56-

Part of the data on each of the buses 193-1 through 193-8 is the operation code which specifies which one of the operations corresponding to the output buses 54-1 through 54-8 is to be selected. That opcode information appears on bus 192 to control the selector 188 to select the desired one of the multiplexers 190-1 through 190-8. With this single operation mode, the input register 53-1 acts as a pipeline for single operations. Up to eight single operations are loaded at one time into the register 53-1. After the single operations are loaded into the register 53-1 over bus 184, each of those operations is selected by multiplexer 187 for output to the appropriate stage of the output register 53-2. Each new instruction loads either multiple operations or single operation information into register 53-1. Each time a multiple operation or a single operation appears on the bus 184, a mode control field appears on line 195 for storage in the mode control register 185.

When the mode control 185 calls for a multiple operation, then the contents of register 53-1 is transferred directly into register 53-2. When the mode control 185 calls for single operation, the operations stored into register 53-1 in parallel are serially unloaded through multiplexers 187 and 188 one at a time, into the output register 53-2.

The computer of system 3, using the instruction unit of FIG. 15, switches readily between multiple operation and single operation modes in order to achieve the most efficient operation of the computer system. For those programs in which single operation execution is efficient, the single operation mode is more desirable in that less address space is required in the instruction cache and other units circuits of the system. For example, up to eight times as many single operation instructions can be stored in the same address space as one multiop instruction. Of course, the number of concurrent multiple

-57-

operations (eight in the FIG. 15 example) is arbitrary and any number of parallel operations for the multiple operation mode can be specified.

Operation

The operation of the invention is described in connection with the execution of a number of programs. Each program is presented in FORTRAN source code form and in the kernel-only code form. The kernel-only code is executed by the computer of FIG. 3 in which the instructions are fetched by the I unit 32-7, the computations are performed by the processors 32 using the data multiconnect 30 and address multiconnect 31.

TABLES 1-1 to 1-6: Kernel-Only Code.

TABLE 1-1 depicts a short vectorizable program containing a DO loop. The loop is executed N times (R=10, see Term Table) for i from 1 to N. The program does not exhibit any recurrence since the result from one iteration of the loop is not utilized in a subsequent iteration of the loop.

-58-

TABLE 1-1

VECTORIZABLE LOOP

```
DO 10 i = 1,N
A = XR(i) * YR(i) - XI(i) * YI(i)
B = XR(i) * YI(i) + XI(i) * YR(i)
XR(i) = A + TR(i)
XI(i) = B + TI(i)
YR(i) = A - TR(i)
YI(i) = B - TI(i)
10  CONTINUE
```

In TABLE 1-2, a listing of the operations utilized for executing the loop of TABLE 1-1 is shown. The operations of TABLE 2-1 correspond to the operations performable by the processors 32 of FIG. 3. Particularly, with reference to FIG. 3, the address add, AAd1, is executed by processor 32-5, address add, AAd2, is executed by processor 32-6, the Mem1read by processor 32-3, the Mem2read by processor 32-4, and the floating-point multiply(FMpy), add(Fadd), subtract(Fsub) by processor 32-1, and the Brtop by the I unit processor 32-7.

In TABLE 1-2, operation 5 by way of example adds the operand @XRI[1], from the address multiconnect 31 (row amc5:36-6) to the operand %r1 from the GPR multiconnect 49 (mc0, column 37-3) and places the result operand @XRI in the address multiconnect 31 (amc5). Operation 9 as another example reads an operand XRI from the Mem2 processor 32-14 and stores the operand in row 4 of the data multiconnect 31 (dmc4). The address from which the operand is accessed is calculated by the displacement adder processor 32-9 which

-59-

adds a literal value of 0(#0 input on line 44 from the instruction) to the operand @XRI from row 5 of the address multiconnect 31 (amc5) which was previously loaded by the result of operation 5. The other operations in TABLE 1-2 are executed in a similar manner.

In TABLE 1-3, the scheduled initial instructions, I_ℓ , for the initial instruction stream IS, where ℓ is 1, 2, ..., 26 are shown for one iteration of vectorizable loop of TABLE 1-1 and TABLE 1-2. In TABLE 1-3, the Iteration Interval(II) is six cycles as indicated by the horizontal-lines after each set of six instructions. Each ℓ^{th} -initial instruction in TABLE 1-3 is formed by a set of zero, one or more operations, $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, initiated concurrently, where $0 \leq n \leq (N-1)$, where N is the number(7 in TABLE 1-3) of concurrent operations and processors for performing operations and where the operation O_n^ℓ is performed by the n^{th} -processor in response to the ℓ^{th} initial instruction. The headings FAdd, FMpy, Mem1, Mem2, AAd1, AAd2 and IU refer to the processors 32-1, 32-2, 32-3, 32-4, 32-5, 32-6, and 32-7, respectively, of FIG. 3. By way of example, instruction 1 uses two operations, AAd1 and AAd2 in processors 32-5 and 32-6. Note that instructions 7,8 and 9 have zero operations and are examples of "NO OP's".

TABLE 1-3 is a loop, LP, of L initial instructions $I_0, I_1, I_2, \dots, I_\ell, \dots, I_{(L-1)}$, where L is 26. The loop, LP, is part of the initial instruction stream IS and in which execution sequences from I_0 toward $I_{(L-1)}$.

In TABLE 1-3, the ℓ designates the instruction number, the OP indicates the number of processors that are active for each instruction.

TABLE 1-3 is divided into J stages($J=5$ for TABLE 1-3), S_n^ℓ , equal to 0, 1, ..., 4. Each operation, $O_n^{k,\ell}$, in an ℓ -instruction has a corresponding stage number.

In TABLE 1-4, the schedule of overlapped instructions is shown for iteration of the vectorizable loop of TABLE

-60-

1-2. In FIG. 1-4, a new iteration of the loop begins for each iteration interval (II) that is, at T1, T7, T13, T19, T25, and so on. The loop iteration that commences at T1 completes at T26 with the Mem1, Mem2 operations. In a similar manner, the loop iteration that commences at T7 completes at T32. The iteration that commences that T13 completes at T38. A comparison of the number of operations, the OP column, between TABLE 1-3 and TABLE 1-4 indicates that the TABLE 1-4 operation on an average includes a greater number of operations per instruction than does the TABLE 1-3. Such operation leads to more efficient utilization of the processors in accordance with the present invention.

-61-

TABLE 1-2

GRAPH REPRESENTATION FOR THE VECTORIZABLE LOOP

```

1      * r1 contains 4
2      *      DO 10 i = 1,N
3      *      A = XR(i) * YR(i) - XI(i) * YI(i)
4      *      B = XR(i) * YI(i) + XI(i) * YR(i)

5      @XRI :- AAd1 ( @XRI[1] %r1 )
6      @XII :- AAd2 ( @XII[1] %r1 )
7      @YRI :- AAd1 ( @YRI[1] %r1 )
8      @YII :- AAd2 ( @YII[1] %r1 )
9      XRI :- Mem2read ( #0 @XRI )
10     XII :- Mem1read ( #0 @XII )
11     YRI :- Mem2read ( #0 @YRI )
12     YII :- Mem1read ( #0 @YII )
13     M1 :- FMpy ( XRI YRI )
14     M2 :- FMpy ( XII YII )
15     M3 :- FMpy ( XII YRI )
16     M4 :- FMpy ( XRI YII )
17     A :- Fsub ( M1 M2 )
18     B :- Fadd ( M3 M4 )
19
20     *      XR(i) = A + TR(i)
21     *      XI(i) = B + TI(i)
22     *      YR(i) = A - TR(i)
23     *      YI(i) = B - TI(i)

24     @TRI :- AAd1 ( @TRI[1] #r1 )
25     @TII :- AAd2 ( @TII[1] #r1 )
26     TRI :- Mem2read ( #0 @TRI )
27     TII :- Mem1read ( #0 @TII )
28     XRIN :- Fadd ( TRI A )
29     XIIN :- Fadd ( TII B )
30     YRIN :- Fsub ( TRI A )
31     YIIN :- Fsub ( TII B )
32     $W1 :- Mem2write ( XRIN #0 @XRI )
33     $W2 :- Mem1write ( XIIN #0 @XII )
34     $W3 :- Mem2write ( YRIN #0 @YRI )
35     $W4 :- Mem1write ( YIIN #0 @YII )

36     $BR :- Brtop ( #loop )

37     *10      CONTINUE

```

-62-

TABLE 1-3

SCHEDULE FOR ONE ITERATION OF THE VECTORIZABLE LOOP

T	OP	FAdd	FMpy	Mem1	Mem2	AA1	AA2	IU
1	2					@YRI	@YRI	
2	2					@XRI	@XII	
3	2					@TRI	@XTII	
4	3			YII	YRI			BR
5	2			XII	XRI			
6	2			TII	TRI			

7	0							
8	0							
9	0							
10	0							
11	1			M4				
12	1			M3				

13	1			M2				
14	1			M1				
15	0							
16	1	B						
17	0							
18	1	A						

19	1	YIIN						
20	1	XIIN						
21	1	YRIN						
22	0							
23	1	XRIN						
24	0							

25	2			W4	W3			
26	2			W2	W1			

-63-

TABLE 1-4

SCHEDULE FOR OVERLAPPED ITERATIONS OF THE VECTORIZABLE LOOP

T	OP	FAdd	FMpy	Mem1	Mem2	AAd1	AAd2	IU
1	2					@YRI	@YII	
2	2					@XRI	@XII	
3	2					@TRI	@TII	
4	3			YII	YRI			BR
5	2			XII	XRI			
6	2			TII	TRI			

9	2					@YRI	@YII	
8	3					@XRI	@XII	
9	3					@TRI	@TII	
10	3			YII	YRI			BR
11	3		M4	XII	XRI			
12	3		M3	TII	TRI			

13	3		M2			@YRI	@YII	
14	3		M1			@XRI	@XII	
15	2					@TRI	@TII	
16	4	B		YII	YRI			BR
17	3		M4	XII	XRI			
18	4	A	M3	TII	TRI			

19	4	YIIN	M2			@YRI	@YII	
20	4	XIIN	M1			@XRI	@XII	
21	3	YRIN				@TRI	@TII	
22	4	B		YII	YRI			BR
23	4	XRIN	M4	XII	XRI			
24	4	A	M3	TII	TRI			

25	6	YIIN	M2	W4	W3	@YRI	@YII	
26	6	XIIN	M1	W2	W1	@XRI	@XII	
27	3	YRIN				@TRI	@TII	
28	4	B		YII	YRI			BR
29	4	XRIN	M4	XII	XRI			
30	4	A	M3	TII	TRI			

31	4	YII	M2	W4	W3			
32	4	XIIN	M1	W2	W1			
33	1	YRIN						
34	1	B						
35	2	XRIN	M4					
36	2	A	M3					

37	4	YIIN	M2	W4	W3			
38	4	XIIN	M1	W2	W1			

-64-

In TABLE 1-5, the kernel-only schedule for the TABLE 1-1 program is shown. The l1 through l6 schedule of TABLE 1-5 is the same as the schedule for the stage including instructions l25 through l30 of TABLE 1-4. The operations of the kernel-only schedule are not all performed during every stage. Each stage has a different number of operations performed. The operations for l1 through l6 are identified as stage A, l7 through l12 are identified as stage B, l13 through l18 are identified as stage C, l19 through l24 are identified as stage D, l25 through l26 are identified as stage E.

TABLE 1-5

KERNEL-ONLY SCHEDULE AND CODE FOR THE VECTORIZABLE LOOP

T	OP	FAdd	FMpy	Mem1	Mem2	AAAd1	AAAd2	IU
1	6	YIIN	M2	W4	W3	@YRI	@YRI	
2	6	XIIN	M1	W2	W1	@YRI	@XII	
3	3	YRIN				@TRI	@TRI	
4	4	B		YII	YRI			BR
5	4	XRIN	M4	XII	XRI			
6	4	A	M3	TII	TRI			

In TABLE 1-6, the operation of TABLE 1-4 overlapped schedule is represented in terms of the stages of TABLE 1-3. For the first iteration($\bar{i}=0$), only those operations of stage A are executed. For the second iteration($\bar{i}=1$), only those operations of stages A and B are executed. For the third iteration($\bar{i}=2$), the operations of stages A, B and C are executed. For the fourth iteration($\bar{i}=3$), the operations of stages A, B, C, and D are executed. Iterations 0, 1, 2, and 3 represent the prolog during which less than all of the operations of the kernel are executed.

The iterations 4, 5, and 6 of TABLE 1-6 represent the body of the loop and all operations of the kernel-only code are executed.

-65-

The iterations 7, 8, 9, and 10 of TABLE 1-6 represent the epilog of the loop and selectively fewer operations of the kernel-only code are executed, namely, B, C, D, and E; C, D, and E; D and E; and E.

The manner in which certain operations of the kernel-only code are enabled and disabled is under control of the predicate values stored in the ICR multiconnect 29. During the first execution, only the operations of the A stage are enabled by a predicate. This condition is indicated in TABLE 1-6 for \bar{i} equal 0 by a 1 in the A column of ICR while all other columns have 0. Similarly, a 1 appears in the A and B columns of the ICR so as to enable the A and B operations in the \bar{i} equal to 1 case, and so on for each case for \bar{i} equal 2 through 10.

-66-

TABLE 1-6

Assume 5 stages, 7 iterations

\bar{i}	lc	esc		ICR				
				A	B	C	D	E
0	6	4	A	1	0	0	0	0
1	5	4	B A	1	1	0	0	0
2	4	4	C B A	1	1	1	0	0
3	3	4	D C B A	1	1	1	1	0
4	2	4	E D C B A	1	1	1	1	1
5	1	4	E D C B A	1	1	1	1	1
6	0	4	E D C B A	1	1	1	1	1
7	0	3	E D C B	0	1	1	1	1
8	0	2	E D C	0	0	1	1	1
9	0	1	E D	0	0	0	1	1
10	0	0	E	0	0	0	0	1

-67-

TABLES 2-1, 2-2, 2-3: Recurrence On Loop.

TABLE 2-1 is an example of a FORTRAN program having a recurrence on a loop. The loop is executed five times ($R=5$) for i from 3 to 7. A recurrence exists because the current value in one iteration of the loop, $F(i)$, is determined using the results, $F(i-1)$ and $F(i-2)$ from previous iterations of the loop.

TABLE 2-1

```

INTEGER F(100)
F(1)=1
F(2)=1
DO 10 i=3,7
F(i)=F(i-1)+F(i-2)
CONTINUE
10      END

```

TABLE 2-2 depicts the initial conditions that are established for the execution of the TABLE 2-1 program using the computer of FIG. 3. Referring to FIG. 3, the GPR multiconnect 49 location 2 is set to 4($\%r2=4$), the data multiconnect 30 row 1, location 63 is set to the value of $F(1)$ set by the program in memory equal to 1 [$dmc\ 1:63=F(1)=1$], the data multiconnect row 1, location 0 is set to the value of $F(2)$ set by the program in memory equal to 1 [$dmc\ 1:0=F(2)=1$], the loop counter 90 of FIG. 4 is set to 4 [$lc=4$], the epilog stage counter is set to 1 [$esc=1$]. With these initial conditions the kernel-only code of TABLE 2-3 is executed to execute the program of TABLE 2-1.

-68-

TABLE 2-2

Required Initial conditions:

```

%r2=4
dmc 1:63=F(1)=1
dmc 1:0=F(2)=1
lc=4
esc=1
amc 5:63 points to F(3)

```

TABLE 2-3

\bar{I}_k	Integer	address	memory
k	<u>IAdd</u>	<u>AAAd1</u>	<u>I unit</u>
			<u>Mem1</u>
0	IAdd{0}	AAAd1{0}	mlwrite{1}
	1:63,1:0	%r2,5:63	1:63,5:1
	=1:62	=5:0	
1		Brtop	
2			
3			

In TABLE 2-3 for each k-instruction, the first line indicates the operation to be performed together with the offset{} of the operation relative to the multiconnect

-69-

pointer, the second line the multiconnect addresses of the source operands, and the third line the multiconnect address of the result operand.

For the 0-instruction, the integer add (processor 32-1 of FIG. 3), IAdd, adds the contents of the data multiconnect 30 location 1:63 (accessed on bus 36-1 from location 63 in row 237-1 of FIG. 3) to the contents of the data multiconnect location 1:0 (accessed on bus 36-2 from location 0 in row 237-1 of FIG. 3) and places the result in data multiconnect location 1:62 (stored on bus 35-1 to location 62 in row 237-1 of FIG. 3) all with an offset of 0 relative to the multiconnect pointer.

For the 0-instruction, the address add1 (processor 32-5 of FIG. 3), AAd1, adds the contents of the GPR multiconnect 49 location %r2 (accessed on bus 36-6 of FIG. 3) to the contents of the address multiconnect location 5:63 (accessed on bus 37-3 from location 63 in row 48-2 of FIG. 3) and places the result in address multiconnect location 5:0 (stored over bus 47-3 to locations 0 in row 48-2 of FIG. 3) all with an offset of 0 relative to the multiconnect pointer. The function of the address add is to calculate the address of each new value of F(i) using the displacement of 4 since the values of F(i) are store in contiguous word address (four bytes).

-71-

TABLE 2-4 depicts the operation of the TABLE 2-3 kernel-only code for the four loops ($R=4$) of TABLE 2-1 and five kernel loops ($\bar{R}=5$). The loop counter is decremented from 4 to 0 and the epilog stage counter is decremented from 1 to 0. The multiconnect address range is 0, 1, 2, 3, ..., 63 and wraps around so that the sequence is 60, 61, 62, 63, 0, 1, 2, ... and so on. The addresses are all calculated relative to the multiconnect pointer (mcp). Therefore, a multiconnect address 1:63 means multiconnect row 1 and location 63+mcp. Since the value of mcp, more precisely $mcp(\bar{i})$, changes each iteration, the actual location in the multiconnect changes for each iteration.

The mcp-relative addressing can be understood, for example, referring to the integer add in TABLE 2-4. The function of the integer add is to calculate $F(i-1)+F(i-2)$ (see TABLE 2-1). The value $F(i-1)$ from the previous iteration ($i-1$) is stored in data multiconnect location 1:63. The value $F(i-2)$ from the 2nd previous iteration ($i-2$) is stored in data multiconnect location 1:0. The result from the add is stored in data multiconnect location 1:62.

Referring to the IAdd instruction at T0, the result is stored at the address $mcp(0)+62$ when $\bar{i}=0$. Referring to the IAdd instruction at T4, one operand is accessed from the address $mcp(1)+63$ when $\bar{i}=1$. However, $mcp(1)$ equals $mcp(0)-1$, so that $[mcp(0)+62]$ equals $[mcp(1)+63]$. Therefore, the operand accessed by the IAdd instruction at T4 is the very same operand stored by the IAdd instruction at T0. Similarly, referring to the IAdd instruction at T8, the operand from 1:0 is accessed from $mcp(2)+0$. However, $mcp(2)$ equals $mcp(0)-2$ and therefore, $[mcp(0)+62]$ equals $[mcp(2)+0]$ and the right-hand operand accessed by IAdd at T8 is the operand stored by IAdd at T0. Note that the execution does not require the copying of result of an operation during one iteration of a loop to another location even though that result is saved for use in

-72-

subsequent iterations and even though the subsequent iterations generate similar results which must also be saved for subsequent use. The invariant addressing using the multiconnect pointer is instrumental in the operations which have a recurrence on the loop.

TABLES 2-1, 2-5, 2-6: Single And Multiple Operations.

TABLES 2-5 and 2-6 represent the single processor execution of the program of TABLE 2-1. TABLE 2-5 represents the initial conditions and TABLE 2-6 represents the kernel-only code for executing the program. In the single operation embodiment, only a single one of the processors of FIG. 3 is utilized during each instruction. Referring to TABLE 2-6 and FIG. 3, the 0-instruction for the IAdd uses processor 32-1, the 1-instruction for the AAd1 uses processor 32-5, the 2-instruction for Brtop uses I unit processor 32-7, and the 3-instruction for mlwrite uses processor 32-3. The invariant and relative addressing of the multiconnect units is the same as in the other examples described. The recurrence on the loop is the same for the TABLE 2-6 example as previously described for TABLE 2-4.

TABLE 2-5

Initial Conditions:

1:63 = F(1) = 1

1:0 = F(2) = 11

5:63 = address in main store
space F(3)

%r2 = 4

-73-

TABLE 2-6

<u>Time</u>	<u>Unit</u>	<u>Operation</u>	<u>Source1</u>	<u>Source2</u>	<u>Destination</u>
0	Integer	IAdd{0}	1:63	1:0	=1:62
1	Address	AAAd1{0}	%r2	5:63	=5:0
2	IU	BRtop			
3	memory	mlwrite{1}	1:63	5:1	

TABLES 3-1 To 3-5: Conditional On Recurrence Path.

TABLE 3-1 is an example of a Fortran program which has a conditional on the recurrence path of a loop. The Fortran program of TABLE 3-1 is executed to find the minimum of a vector. The vector is the vector X which has a hundred values. The trial minimum is XM. The value of XM after completing the execution of the loop is the minimum. The integer M is the trial index and the integer K is the loop index.

The initial conditions for the kernel-only code for executing the loop of TABLE 3-1 are shown in TABLE 3-2. The general purpose register offset location 1 is set equal to 1 so that the value of K in TABLE 3-1 can be incremented by 1. The general purpose register offset location 4 is set equal to 4 because word (four bytes) addressing is employed. The addresses of the vector values are contiguous at word location. The multiconnect temporary value ax[1] is set equal to the address of the first vector value X(1). Similarly, the multiconnect temporary location XM[1] is set equal to the value of the first vector value X(1).

-74-

TABLE 3-1

EXAMPLE WITH CONDITIONAL ON RECURRENCE PATH
FIND THE FIRST MINIMUM OF A VECTOR

```
INTEGER M,K  
REAL X(100), XM  
  
XM = X(1)  
M = 1  
DO 10 K = 2,100  
IF (X(K).GT.XM) GO TO 10  
XM = X(K)  
M = K  
10 CONTINUE
```

TABLE 3-2

Initial Conditions:

```
%r1=1, %r4=4  
ax[1] = points to addr(x(1))  
m[1] = x(1)  
xm[1] = has value x(1)
```

In TABLE 3-3, the kernel-only code for the program of
TABLE 3-1 is shown.

-75-

In TABLE 3-1, the recurrence results because the IF statement uses the trial minimum XM for purposes of comparison with the current value X(K). However, the value of XM in one iteration of the loop for the IF statement uses the result of a previous execution which can determine the value of XM as being equal to X(K) under certain conditions. The conditions which cause the determination of whether or not XM is equal to X(K) are a function of the Boolean result of comparison "greater than" in the IF statement. Accordingly, the TABLE 3-1 program has a conditional operation occurring on the recurrence path of a loop.

TABLE 3-3 is a representation of the kernel-only code for executing the program of TABLE 3-1. Referring to TABLE 3-3, the recurrence exists and is explained by the follow cycle of operand dependency within the kernel-only code. For convenience reference is first made to the "greater than" comparison, Fgt, which occurs in instruction 11. The comparison in instruction 11 is between the trial minimum XM and the current value being tested, X, as derived from the vector. The comparison in instruction 11 creates a Boolean result and that result is then used in the next iteration of the loop when it is stored into a predicate by the Stuffbar operation of instruction 3. In instruction 3, the Boolean value is stored into the predicate at the sw [2] location, and that location is thereafter used in instruction 6 to select the new trial minimum using the FMsel operation. The index for the trial minimum is stored in instruction 6 using the Isel operation. The trial minimum XM stored in instruction 6 by the FMsel operation is then used again in the instruction 11 to do the "greater than" comparison, thereby returning to the starting point of this analysis.

-76-

TABLE 3-3

\bar{I}_k k	Alu/ IAdd	FMpy	Memory Mem1	Memory Mem2	AAd1	IU
0	IAdd{0} k[1],%r1 =k				AAd1{0} ax[1],%r4 =ax	
1						
2						
3			Stuffbar{2} Isel[2] =sw[2]	m2read{2} ax =x		
4						
5						
6	Isel{sw[2]} k[2],m[3] =m[2]	FMsel{sw[2]} x[2],xm[3] =xm[2]				
7						
8						
9						
10						Brtop
11	Fgt{1} x[1],xm[2] =Isel[1]					

-77-

The TABLE 3-2 and TABLE 3-3 example above utilized variable names for the locations within the multiconnect units.

TABLE 3-4 and TABLE 3-5 depict the same example using absolute multiconnect addresses (relative to mcp).

TABLE 3-4

Initial Conditions:

```
%r1=1, %r4=4
5:1 <= points to Addr(x(1))
1:11 <= x(1)(DMC)
2:1 <= x(1)
```

-78-

TABLE 3-5

<u>\bar{I}_k</u>	<u>Alu/IAdd</u>	<u>FMpy</u>	<u>Mem1</u>	<u>Mem2</u>	<u>AAAd1</u>	<u>IU</u>
0	IAdd{0} 1:1,%r1 =1:0				AAAd1{0} 5:1,%r4 =5:0	
1						
2						
3			Stuffbar{2} 1:22 =icr12	m2read{0} 5:0 =4:0		
4						
5						
6	Isel{12} 1:2,1,11 =1:12	FMsel{12} 4:2,2:2 =2:2				
7						
8						
9						
10						Brtop
11	Fgt{1} 4:1,2:2 =1:21					

TABLES 4-1 TO 4-4: Complex Conditionals Within Loop

The TABLE 4-1 is an example of a complex program having a conditional on the loop, but without any recurrence. Of course, recurrence processing can be handled in the same manner as the previous examples.

TABLE 4-2 depicts the kernel-only code with the initial conditions indicated at the top of the table.

TABLE 4-3 depicts the multiconnect addresses (relative to mcp) for the TABLE 4-2.

TABLE 4-1

```
DO 12 I=1,N
1  X = D(I)
2  IF (X.LT.4) GO TO 9
3  X = X + 3
4  X = SQRT(X)
5  IF (X.LT.3) GO TO 12
6  X = X + 4
7  X = SQRT(X)
8  GO TO 12
9  IF (X.GT.0) GO TO 5
10 X = X + 5
11 X = 2*x
12 D(I) = X
```

-80-

TABLE 4-2

%r1=1, %r4=4

ax[1] <- points to addr(x(0)) element prior to x(1)

	FAdd/lalu	Multiplier	Meml	AAddl	IU
0	Or{3} j95[3],j45[3] =j5[3]			AAdd {0} ax[1],%r0 =ax[0]	*
1	Isel{d3[3]} x4[3],x1[3] =x5i[3]		Stuffbar{2} c2[2] =d3[2]		*
2	FAdd{4} x5i[4],%r4 =x6[4]		Stufficr{2} c2[2] =d9[2]		
3	Fge{2} x1[2],%r2 =nc2[2]	FMPy{d10[3]} x10[3],%r2 =x11[3]	miread{0} ax[0],%r0 =x1[0]		
4	FAdd{d3[2]} x1[2],%r3 =x3[2]		Stufficr{3} j5[3] =d5[3]		
5	Fgt{d9[2]} x1[2],%r0 =c9[2]				
6	Isel{d10[5]} x11[5] is12p[5] =is12[5]	Sqrt{d6[4]} x6[4] =x7[4]			
7	Flt{d6[3]} x5i[3],%r4 =c5[3]				
8	Isel{2} %r1,%r0 =j45[2]	Sqrt{d3[2]} x3[2] =x4[2]			
9	Isel{d9[2]} %r1,%r0 =j95[2]		Stuffbar{d9[2]} c9[2] =d10[2]		
10	Flt{1} x1[1],%r4 =c2[1]		mwrite{5} is12[5]		Brtop
11	Isel{d6[4]} x7[4],x4[4] =is12p[4]		Stuffbar{d5[3]} c5[3] =d6[3]		
12	FAdd{d10[2]} x1,%r5 =x10[2]				

-81-

TABLE 4-3

Register Assignments:

j5[3] = 1:12
ax[0] = a1:0
x5i[3] = 1:6
d3[2] = icr:12
x6[4] = 1:0
d9[2] = icr:14
nc2[2] = 1:2
x11[3] = 2:0
x1[0] = 3:0
x3[2] = 1:11
d5[3] = icr:11
c9[2] = 1:10
is12[5] = 1:15
x7[4] = 2:6
c5[3] = 1:14
j45[2] = 1:8
x4[2] = 2:3
j95[2] = 1:7
c2 = 1:9
is10[4] = 1:0
d6[3] = icr9
x10[2] = 1:12
is12p[4] = 1:30

-82-

While the invention has been particularly shown and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and details may be made therein without departing from the spirit and scope of the invention.

-83-

What is claimed is:

1. A computer system including a processing unit having one or more processors for performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit.

2. A computer system employing a horizontal architecture comprising:

instruction processing means for controlling the execution of an instruction stream where the instruction stream includes instructions specifying operations to be performed and for providing address offsets for use in connection with the operations to be performed,

one or more processors for performing operations specified by said instructions, where each processor includes,

one or more functional units for forming output operands from input operands,

one or more processor input ports for providing input operands to said functional units,

one or more processor output ports for providing output operands from said functional units,

-84-

a plurality of multiconnect elements, addressed by multiconnect addresses, for use in providing operands from and to said processors where each multiconnect element includes,

one or more multiconnect cells for storing operands,

multiconnect access means for accessing one or more of said multiconnect cells addressed by a multiconnect address,

one or more multiconnect input ports for providing output operands to accessed ones of said multiconnect cells,

one or more multiconnect output ports for providing input operands from accessed ones of said multiconnect cells,

processor-multiconnect interconnection means for connecting output operands from processor output ports to multiconnect input ports,

multiconnect-processor interconnection means for connecting input operands from multiconnect output ports to processor input ports,

invariant addressing means including,

pointer register means for storing multiconnect pointers

address generation means for repeatedly combining pointers with offsets to form said multiconnect addresses,

modifying means for modifying said pointers to form modified pointers whereby said address generation means combines modified pointers and offsets to form modified multiconnect addresses.

3. The system of Claim 2 wherein,
said instruction stream includes a loop of instructions in which execution of at least some of the instructions in the loop is repeated a plurality of times,

-85-

said modifying means modifies pointers in response to repetitions of said loop.

4. The system of Claim 2 wherein said instruction stream includes a loop of instructions in which execution of the loop is repeated a plurality of times, and wherein said instruction stream includes within said loop an instruction providing an address offset,

and wherein said modifying means includes means to modify said pointer in response to repetitions of said loop to form a modified pointer,

whereby the multiconnect element address for said instruction is determined for one iteration of said loop as a function of said address offset and said pointer and is determined for another iteration of said loop as a function of said address offset and said modified pointer.

5. The system of Claim 3 wherein said pointer is $mcp(\bar{i})$ and wherein said modifying means includes means to modify said pointers, $mcp(\bar{i})$, in response to repetitions of said loop to form modified pointers, $mcp(\bar{i}+1)$ equal to $mcp(\bar{i})^*$,

whereby the multiconnect address for said instruction is determined for one iteration of said loop as a function of $mcp(\bar{i})$ and said address offset and is determined for another iteration of said loop as a function of said address offset and $mcp(\bar{i}+1)$ which is equal to $mcp(\bar{i})^*$.

6. The system of Claim 5 wherein

said address generation means includes an adder for adding said pointer with offsets, ao , to form said multiconnect addresses, and said modifying means includes means to decrement said pointer, $mcp(\bar{i})$, by one for each iteration of said loop to form said modified pointers, $mcp(\bar{i}+1)$ equal to $mcp(\bar{i})-1$ equal to $mcp(\bar{i})^*$, $mcp(\bar{i}+2)$ equal to $mcp(\bar{i})-2$ equal to $mcp(\bar{i}+1)^*$, $mcp(\bar{i}+3)$ equal to $mcp(\bar{i})-3$

-86-

equal to $mcp(\bar{i}+2)^*$, $mcp(\bar{i}+n)$ equal to $mcp(\bar{i})-n$ equal to $mcp(\bar{i}+n-1)^*$,

whereby the multiconnect addresses for said instruction are determined for n-iterations of said loop as a function of said offsets, ao, and respectively, $mcp(\bar{i}+1)$, $mcp(\bar{i}+2)$, $mcp(\bar{i}+3)$, ..., $mcp(\bar{i}+n)$ to form multiconnect addresses $ao+mcp(\bar{i}+1)$, $ao+mcp(\bar{i}+2)$, $ao+mcp(\bar{i}+3)$, ..., $ao+mcp(\bar{i}+n)$, respectively.

7. The system of Claim 2 where said address generation means combines a first offset from an instruction with a first pointer to form a first multiconnect address at one time and combines a second offset from an instruction with a modified first pointer to form said first multiconnect address at a second time whereby said multiconnect element is accessed at the same location at different times.

8. The system of Claim 2 where said address generation means combines a first offset from an instruction with a first pointer to form a first multiconnect address at one time and combines said first offset with a modified first pointer to form a second multiconnect address at a second time whereby said multiconnect element is accessed at different locations at different times using the same offset.

9. The system of Claim 2 where the processor-multiconnect interconnection means connects a first set of multiconnect elements with a first processor output such that each output operand from said processor output can be written into one or more of said first set of multiconnect elements.

10. The system of Claim 2 where the multiconnect-processor interconnection means connects a first processor input to a second set of multiconnect elements such that

-87-

any input operand within said second set of multiconnect elements can be read into said processor input.

11. The system of Claim 2 where,

said processor-multiconnect interconnection means connects a first set of multiconnect elements with a first processor output such that each output operand from said processor outputs can be written into one or more of said first set of multiconnect elements,

said multiconnect-processor interconnection means connects a first processor input to a second set of multiconnect elements such that any input operand residing within said second set of multiconnect elements can be read into said processor input,

said first set of multiconnect elements and said second set of multiconnect elements each including one multiconnect element in common whereby the first processor output operand can form the first processor input operand.

12. The system of Claim 2 where,

said processor-multiconnect interconnection means connects, for each processor, a processor output to a first set of multiconnect elements such that each output operand from said processor output can be written into one or more of said first set of multiconnect elements,

said multiconnect-processor interconnection means connects, for each processor, a processor input from a second set of multiconnect elements such that any input operand residing within said second set of multiconnect elements can be read into said processor input,

said first set of multiconnect elements and said second set of multiconnect elements, for each processor, including one multiconnect element in common whereby the processor output operand from any processor can form the processor input operand for any processor.

-88-

13. The system of Claim 12 where,
each output operand, for each processor, is transmitted from the processor output into every multiconnect element in the first set of multiconnect elements associated with each said processor.

14. The system of Claim 12 where,
each processor can simultaneously on each cycle read from a multiconnect element one input operand, using a multiconnect element output port and a processor input port, and write one output operand to a multiconnect element, using a processor output port and a multiconnect element input port.

15. The system of Claim 12 wherein said first pointer is used in combination with said first offset address to store an operand at a first multiconnect element address in a multiconnect element at one time and said first pointer after modification is again used in combination with said second offset address to read said operand from said first multiconnect element address.

16. The system of Claim 15 wherein there is a different pointer for each multiconnect element.

17. The system of Claim 16 wherein all of said pointers are equal.

18. A data processing system employing a horizontal architecture comprising:

instruction processing means for controlling the execution of an instruction stream where the instruction stream includes a loop of instructions in which execution of at least some of the instructions in the loop is repeated a plurality of times,

-89-

one or more processors for use in connection with the execution of said instructions, where each processor includes,

- one or more functional units for forming output operands from input operands,
- one or more processor input ports for inputting input operands to functional units,
- one or more processor output ports for outputting output operands from said functional units,

a plurality of connection multiconnect elements including first and second multiconnect elements addressed by multiconnect element addresses where each multiconnect element includes,

- one or more multiconnect cells for storing operands,

- multiconnect element access means for accessing one or more of said multiconnect cells addressed by a multiconnect element address,

- one or more multiconnect element input ports for inputting output operands to accessed ones of said multiconnect cells,

- one or more multiconnect element output ports for outputting operands from accessed ones of said multiconnect cells,

processor-multiconnect interconnection means for connecting output operands from processor output ports to multiconnect element input ports,

multiconnect-processor interconnection means for connecting input operands from multiconnect element output ports to processor input ports,

invariant addressing means including,

- pointer register means for storing pointers, including first pointer and second

-90-

pointers associated with said first multiconnect element and a second-memory pointer associated with said second multiconnect element,

offset register means for storing offset addresses, including first and second offset addresses associated with said first multiconnect element and a second-memory offset address associated with said second multiconnect element,

address generation means for combining pointers and offset addresses to form multiconnect element addresses, said first pointer and first offset addresses for said first-memory being combined to form a first-memory address with a first particular value at one time, and said second pointer and second offset addresses for said first-memory being combined to form said first-multiconnect element address with a second particular value at another time wherein said first and second particular values are equal,

means for modifying said pointers in sequence in response to repetitions of said loop of instructions such that said first and second particular values change but remain equal whereby said first and second particular values after each modification of pointers provide a first-multiconnect address for addressing the same multiconnect element cell.

19. A computer system employing horizontal architecture comprising:

a plurality of processors, each processor having processing input terminals and processing output terminals, said processors operable concurrently and in parallel,

source address means for providing source addresses for data to be sourced to each of said processors,

-91-

destination address means for providing destination addresses, for data from each of said processors,

an interconnect circuit for coupling said processing output terminals to said processing input terminals to provide data paths among said processors, said interconnect circuit including,

a plurality of multiconnect elements logically arranged in rows and columns, each of said multiconnect elements having multiconnect element locations for storing information, multiconnect element address means for addressing said locations, multiconnect element data input terminals for connecting information into said multiconnect element locations, multiconnect element data output terminals for connecting information from said multiconnect element locations, multiconnect element row address terminals for connecting row addresses to said multiconnect element address means, and multiconnect element column address terminals for connecting column addresses to said multiconnect element address means,

a plurality of row means, each row means including row data means connecting all of the multiconnect element data input terminals in a row in common to the processing output terminals of a corresponding processor, each row means including row address means connecting all of the multiconnect element address input terminals in a row in common to the destination address means whereby all multiconnect elements in a row store the same information from said corresponding processor at the same row address location,

a plurality of column means, each column means including column address means connecting all of the column address terminals in a column in common to the source address means, each column means including column data means connecting all of the multiconnect

-92-

element data output terminals in a column in common to the processing input terminals of a corresponding processor whereby each processor can receive information from the output of any other processor,

sequence control means for providing said source and destination addresses for each processor whereby said multiconnect elements are addressed in common.

20. The system of Claim 2 wherein each multiconnect element includes,

a multiconnect element pointer register for storing an offset address,

an adder for adding said offset address to said source address to form a read address for addressing said locations.

21. The system of Claim 2 wherein each multiconnect element includes,

a multiconnect element pointer register for storing an offset address,

an adder for adding said offset address to said source address to form a read address for addressing said locations,

an incrementer for changing said multiconnect element pointer register to change said offset.

22. The system of Claim 7 wherein said incrementer operates to decrement said multiconnect element pointer register by 1 to change said offset.

23. A computer system employing a horizontal architecture for use with an instruction stream where the instruction stream includes a number of instructions, \bar{I}_0 , \bar{I}_1 , \bar{I}_2 , ..., \bar{I}_k , ..., $\bar{I}_{(K-1)}$ of an instruction stream, \bar{IS} , where each said instruction, \bar{I}_k , of the instruction stream

-93-

\bar{IS} specifies operations $\bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_N^{k,\ell}$, where each operation, $\bar{O}_n^{k,\ell}$, provides address offsets, $ao_n^k(c)$, comprising,

instruction processing means for sequentially accessing said instructions, \bar{I}_k , and corresponding operations, $\bar{O}_n^{k,\ell}$, said instruction processing means accessing said instructions one or more times during one or more iterations, \bar{i} , of said instruction stream \bar{IS} ,

one or more processors, each processor for performing one or more of said operations specified by said instructions, \bar{I}_k , said processors including input and output ports,

a plurality of multiconnect memories, addressed by memory addresses, $a_n^k(c)(\bar{i})$, for connecting operands from and to said processors, said memories having input and output ports, said memories providing input operands on said memory output ports when addressed by said memory addresses,

processor-memory interconnection means for connecting output operands from processor output ports to memory input ports,

memory-processor interconnection means for connecting input operands from memory output ports to processor input ports,

invariant addressing means for addressing said memories during different iterations including a current iteration, \bar{i} , and a previous iteration, $(\bar{i}-1)$ including,

modifying means for forming a current pointer address, $mcp(\bar{i})$, from a previous pointer address $mcp(\bar{i}-1)$ with the operation $D^*[mcp(\bar{i}-1)]$ such that $mcp(\bar{i})=D^*[mcp(\bar{i}-1)]$,

pointer register means for storing said pointer address, $mcp(\bar{i})$, for use in the \bar{i}^{th} -iteration,

address generation means for combining the pointer address, $mcp(\bar{i})$, with an address offset,

-94-

$ao_n^k(c)$, to form said memory addresses, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration.

24. The system of Claim 22 wherein, said modifying means includes means for forming said current pointer address as a new value in response to each iteration, \bar{i} .

25. The system of Claim 23 wherein said modifying means includes means for performing said operation, $D^*[]$, by adding or subtracting a number to or from $mcp(\bar{i}-1)$ to form $mcp\bar{i}$.

26. The system of Claim 25 wherein said number is relatively prime to the maximum value of said pointer address.

27. A computer system employing a horizontal architecture for use with an initial instruction stream, IS, where the initial instruction stream includes a loop, LP, of initial instructions, including initial instructions $I_0, I_1, I_2, \dots, I_\ell, \dots, I_{(L-1)}$ in which execution commences with I_0 one or more times, once for each iteration, i , of the loop, LP, where each said initial instruction, I_ℓ , specifies operations to be performed, including the initial operations $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, where said initial instructions have been transformed to a kernel, KE, of kernel instructions including the kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ in which execution commences with \bar{I}_0 one or more times, once for each iteration, \bar{i} , of the kernel, KE, where each said kernel instruction, \bar{I}_k , specifies kernel operations to be performed, including the kernel operations $\bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_N^{k,-}$, where each kernel operation, $\bar{O}_n^{k,\ell}$, provides address offsets, $ao_n^k(c)$, for use in connection with the kernel operations to

-95-

be performed, and where instruction execution is measured by instruction periods, T , comprising,

instruction processing means for sequentially accessing said kernel instructions, \bar{I}_k , and corresponding kernel operations, $\bar{O}_n^{k,\ell}$, said instruction processing means accessing said instructions one or more times during one or more iterations, \bar{I} , of said instruction stream $\bar{I}S$,

one or more processors, each processor for performing one or more of said kernel operations specified by said instructions, \bar{I}_k , each processor having one or more input ports and output ports,

a plurality of multiconnect memories, addressed by memory addresses, $a_n^k(c)(\bar{I})$, for connecting operands from and to said processors, said memories having input and output ports, said memories providing input operands on said memory output ports when addressed by said memory addresses,

processor-memory interconnection means for connecting output operands from processor output ports to memory input ports,

memory-processor interconnection means for connecting input operands from memory output ports to processor input ports,

invariant addressing means for addressing said memories during different iterations including a current iteration, \bar{I} , and a previous iteration, $(\bar{I}-1)$ including,

modifying means for forming a current pointer address, $mcp(\bar{I})$, from a previous pointer address $mcp(\bar{I}-1)$ with the operation $D^*[mcp(\bar{I}-1)]$ such that $mcp(\bar{I})=D^*[mcp(\bar{I}-1)]$,

pointer register means for storing said pointer address, $mcp(\bar{I})$, for use in the \bar{I}^{th} -iteration,

address generation means for combining the pointer address, $mcp(\bar{I})$, with an address offset,

-96-

$ao_n^k(c)$, to form said memory addresses, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration.

26. The system of Claim 27 wherein execution of the kernel is repeated a plurality of times including said $(\bar{i}-1)^{th}$ -iteration and said \bar{i}^{th} -iteration and wherein executions of the kernel causes the multiconnect pointer to change, wherein said kernel includes an instruction providing an address offset, ao_n^k , and wherein said modifying means modifies said previous pointer addresses, $mcp(\bar{i}-1)$, in response to repetitions of said kernel to form said current pointer addresses, $mcp(\bar{i})$,

whereby said address generation means forms the memory address, $a_n^k(c)(\bar{i}-1)$, for said instruction, \bar{I}_k , for the $(\bar{i}-1)^{th}$ -iteration of said kernel as a function of said address offset ao_n^k and the previous pointer address $mcp(\bar{i}-1)$ and forms the memory address, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration of said kernel as a function of said address offset $ao_n^k(c)$ and the current pointer address $mcp(\bar{i})$.

29. The system of Claim 28 wherein execution of the kernel occurs for iterations 1, 2, ..., $(\bar{i}-1)$, \bar{i} , ..., \bar{R} ,

said address generation means includes an adder for adding said current pointer address, $mcp(\bar{i})$, with offsets, ao_n^k , to form said memory addresses, $a_n^k(c)(\bar{i})$, and

said modifying means includes means to add one to said previous pointer address, $mcp(\bar{i}-1)$, for each iteration of said kernel to form said current pointer address equal to $mcp(1)$, $mcp(2)$, $mcp(3)$, ..., $mcp(\bar{i}-1)$, $mcp(\bar{i})$, ..., $mcp(2)$,

whereby the memory addresses for said kernel instruction are determined for said kernel to form $ao_n^k(c)+mcp(1)$, $ao_n^k(c)+mcp(2)$, $ao_n^k(c)+mcp(3)$, ..., $ao_n^k(c) + mcp(\bar{i}-1)$, $ao_n^k(c) + mcp(i)$, ..., $ao_n^k(c)+mcp(X)$, respectively.

30. The system of Claim 27 wherein said instruction processing means includes,

-97-

a loop counter for storing a loop count, "lc" representing the number of iterations of said loop, LP, said loop count having a loop count range including a loop count end count,

an epilog counter for storing an epilog count, "esc" representing the ending iterations of said loop, KE, said epilog count having an epilog count range including an epilog end count,

counter control means for controlling the loop counter and epilog counter.

31. A computer system as set forth in Claim 30, wherein said instruction processing means further includes:

counter control means for controlling "lc" and "esc" in accord with the following operations,

if the "lc" is negative, the "esc" is negative, or if the "lc" and "esc" are both zero, then the branch is not taken; otherwise, the branch is taken,

if the "lc" is greater than zero, then the "lc" is decremented; otherwise, it is unchanged,

if the "lc" is zero, and the "esc" is greater than or equal to zero, then the "esc" is decremented; otherwise, it is unchanged,

if the "lc" is positive, and the "esc" is greater than or equal to zero, then the "mcp" is decremented.

32. A computer system as set forth in Claim 27 wherein said instruction processing means includes,

means for storing an iteration value representing different iterations of said kernel, KE,

processor control means for enabling different ones of said processors as a function of said iteration value whereby different ones of the kernel operations from the same kernel instructions are executed during different iterations.

-98-

33. A computer system as set forth in Claim 32 herein said kernel, KE, and the kernel instructions, $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_K$, are iteratively executed during sequential executions of the kernel loop, KE, and wherein operations in each of the kernel instructions are selectively enabled during different iterations of the kernel loop, KE, under control of said processor control means.

34. A computer system as set forth in Claim 33 wherein said processor control means includes iteration control register means for storing control information for selectively enabling the operations, $\bar{O}_n^{k,l}$ during repetitive executions of the kernel, KE, and includes operation address means for addressing said iteration control register means.

35. A computer system as set forth in Claim 32 wherein said one or more processors include n processors, one for each of said n operations to be performed by said instructions, \bar{I}_k , said system further including iteration control means for enabling said processors selectively during different iterations.

-99-

36. A computer system as set forth in Claim 35 wherein said iteration control means includes for each n^{th} -processor,

an iteration control register for storing iteration control signals,

means for storing an iteration control pointer $\text{icp}(\bar{i})$,

means for providing a predicate offset $\text{po}_n^{k,\ell}$ for the $\bar{O}_n^{k,\ell}$ operation to be executed by the n^{th} -processor,

means for combining said predicate offset $\text{po}_n^{k,\ell}$ with said iteration pointer $\text{icp}(\bar{i})$ to form an iteration control address $\text{p}_n^k(\bar{i})$ for addressing said iteration control register to provide an iteration control signal, $C_n(\bar{i})$, for controlling the enabling of said n^{th} -processor during the \bar{i}^{th} -iteration.

37. A computer system as set forth in Claim 36 including N of said processors with n having values 1 to N, wherein said instruction \bar{I}_k provides up to N operations $\bar{O}_n^{k,\ell}$ during each instruction cycle, wherein during each instruction cycle up to N iteration control bits, C_1, C_2, \dots, C_N , are provided, one for each of said operations, respectively, for controlling the enabling of said processors, respectively during said cycle.

38. A computer system as set forth in Claim 37 wherein said instruction unit commences execution of operations for a different one of said kernel instructions, \bar{I}_k , forming said kernel, KE, of K instructions $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_K$ once each instruction period T repetitively over the iteration interval, II, from 1 to K, and wherein said multiconnect pointer is modified to a new value once per iteration interval.

-100-

39. A computer system as set forth in Claim 38 wherein execution of said kernel instructions for multiple iterations of said loop, LP, occur during a prolog period, a kernel period and an epilog period, wherein,

during said prolog period said control bits enable increasing numbers of said processors in successive cycles so that an increasing number of operations become executed,

during said kernel period said control bits enable a maximum number of said processors during iteration interval cycles so that a maximum number of operations become executed,

during said epilog period said control bits enable decreasing numbers of said processors in successive cycles so that a decreasing number of operations become executed.

40. The system of Claim 39 wherein said invariant addressing means includes,

a loop counter for storing a count, "lc" representing the number of iterations of said loop, LP,

an epilog counter for storing a count, "esc" representing the iterations of said loop, LP, during said epilog period,

counter control means for controlling "lc" in a loop counter and "esc" in an epilog counter.

41. A computer system as set forth in Claim 40, wherein said counter control means includes means for controlling "lc" and "esc" in accordance with the following operations,

if "lc" is negative or "esc" is negative, or if "lc" and "esc" are both zero, then the branch is not taken; otherwise, the branch is taken,

if "lc" is greater than zero, then "lc" is decremented; otherwise, it is unchanged,

-101-

if "lc" is zero, and "esc" is greater than or equal to zero, then "esc" is decremented; otherwise, it is unchanged,

if "lc" is positive, and "esc" is greater than or equal to zero, then $mcp(\bar{i})$ is modified.

42. A computer system as set forth in Claim 30, wherein said instruction processing means further includes for each new value of \bar{i} :

means for detecting if "lc" is not equal to the iteration end count,

means for stepping "lc" toward said loop end count if lc is not equal to the loop end count,

means for detecting if "esc" is not equal to the epilog end count,

means for stepping "esc" toward said epilog end count if "esc" is not equal to the epilog end count and "lc" is equal to the loop end count.

43. The system of Claim 8, including kernel loop control means for controlling \bar{R} iterations of said kernel loop, KE,

said kernel loop control means including counter means storing \bar{i} for counting each iteration of said kernel loop, KE,

detector means for detecting when said counter means has been stepped \bar{R} counts whereby further iterations of said kernel loop are terminated.

44. The system of Claim 33 wherein said counter means includes a loop counter for counting R iterations of said kernel loop, and includes an epilog stage counter for counting (S1) iterations of said kernel loop whereby \bar{R} equals $R+(S-1)$.

-102-

45. The computer system of Claim 27 wherein said instruction processing means includes,

counting means for counting iterations, \bar{i} , for the kernel loop, KE, over a count range \bar{R} , and

branch means for exiting execution of said kernel loop after \bar{R} iterations of \bar{i} .

46. The computer system of Claim 45 wherein said counter means includes,

prolog means for counting iterations of \bar{i} during a prolog having the first S-1 iterations of \bar{i} within the count range, \bar{R} ,

body means for counting iterations of \bar{i} during a body having the next R-2 (S-1) iterations of \bar{i} within the count range \bar{R} ,

epilog means for counting iterations of \bar{i} during an epilog having the last S-1 iterations of \bar{i} within the count range \bar{R} .

47. The computer system of Claim 27 wherein said instruction processing means includes operation control means for selecting different ones of the operations to be performed during each iteration \bar{i} of the kernel loop, KE and wherein said kernel loop includes a prolog, body and epilog.

48. The computer system of Claim 47 wherein said operation control means includes,

prolog control means for selecting an increasing number of operations to be performed for each successive iteration of \bar{i} during the prolog.

49. The computer system of Claim 47 wherein said operation control means includes,

body control means for selecting a constant number of operations to be performed during the body.

-103-

50. The computer system of Claim 47 wherein said operation control means includes,

epilog control means for selecting a decreasing number of operations to be performed for each successive iteration of the epilog.

51. The computer system of Claim 27 wherein said counting means includes a loop counter for counting iterations \bar{i} over the count range R and includes an epilog stage counter for counting over the count range S-1 whereby said loop counter and said epilog stage counter together count over the count range \bar{R} equal to $R-(S-1)$.

52. The computer system of Claim 27 wherein each of said processors is associated with iteration control means to enable said processor in response to said instruction processing means during each selected iteration \bar{i} and wherein said instruction processing means includes means for selectively setting said iteration control means for enabling one or more operation in one or more of said processors during each iteration, \bar{i} .

53. The computer system of Claim 52 wherein said iteration control means is operative to control the operation in an enabled processor for each iteration \bar{i} of the kernel loop, KE.

54. The computer system of Claim 53 wherein said iteration control means includes an iteration control register for storing values for controlling the enabling of said processors and operations during each iteration of the kernel loop, and includes means for loading said iteration control register with different iteration control values for each iteration \bar{i} of the kernel loop, KE.

-104-

55. The computer system of Claim 54 including means for loading said iteration control register for each iteration i of the kernel loop, KE, to enable an increasing number of operations to be performed during each successive iteration of the prolog of the kernel loop, KE.

56. The computer system of Claim 54 wherein said means for loading the iteration control register includes means for enabling the same number of operations during each successive iteration of the body of the kernel loop, KE.

57. The computer system of Claim 54 wherein said means for loading the iteration control register includes means to enable a decreasing number of operations during each successive iteration of the epilog of the kernel loop, KE.

58. The computer system of Claim 27 wherein said multiconnect memories includes means for storing results of each operation for each iteration in unique locations for each iteration.

59. The computer system of Claim 58 wherein said instruction processing means includes means for specifying the source for each operation of each iteration.

60. The computer system Claim 59 wherein said instruction processing means includes means for specifying the results for one or more previous iterations as a source for processing in the current iteration.

61. A computer system for performing one or more iterations of a loop of instructions, including, a processing unit having one or more processors for

-105-

performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying for each instruction and each iteration of said loop of instructions operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit,

iteration control means for controlling which operations are active in each instruction during each iteration of said loop.

62. The computer of Claim 61, wherein said iteration control means includes,

an iteration control register formed as a row of multiconnect elements, where one or more of said processors is connected to a corresponding multiconnect element,

processor means for controlling the storing of data into said iteration control register under control of instructions from said instruction unit, wherein said processors each receive an input from the iteration control register for controlling the operation of said processor in response to the data value of the control information from said instruction control register.

-106-

63. A computer system for performing one or more iterations of a loop of instructions, where a result operand from one iteration of a loop is used as a source operand in a subsequent iteration of said loop, including,

a processing unit having one or more processors for performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying for each instruction and each iteration of said loop of instructions operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit.

64. The system of Claim 63 wherein said instruction unit includes,

a loop counter for storing a loop count, "lc" representing the number of iterations of said loop, LP, said loop count having a loop count range including a loop count end count,

an epilog counter for storing an epilog count, "esc" representing the ending iterations of said loop, KE, said epilog count having an epilog count range including an epilog end count,

counter control means for controlling the loop counter and epilog counter.

-107-

65. A computer system as set forth in Claim 63 wherein said instruction unit includes,

means for storing an iteration value representing different iterations of said kernel, KE,

processor control means for enabling different ones of said processors as a function of said iteration value whereby different ones of the kernel operations from the same kernel instructions are executed during different iterations.

66. A computer system for performing one or more iterations of a loop of instructions, where a result operand from one iteration of a loop is used as a source operand in a subsequent iteration of said loop, including,

a processing unit having one or more processors for performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying for each instruction and each iteration of said loop of instructions operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit, wherein said invariant addressing means stores the result operand at a first multiconnect address specified by a first address offset and a first value of said modifiable pointer during one iteration, and wherein said invariant addressing means accesses said result operand as a source operand from said multiconnect address specified by a combination of said

-108-

second address offset and a second value of said modifiable pointer during another iteration of said loop,

iteration control means for controlling which operations are active in each instruction during each iteration of said loop.

67. A computer system employing a horizontal architecture for use with an initial instruction stream, IS, where the initial instruction stream includes a loop, LP, of initial instructions, including initial instructions $I_0, I_1, I_2, \dots, I_\ell, \dots, I_{(L-1)}$ in which execution commences with I_0 one or more times, once for each iteration, i , of the loop, LP, where each said initial instruction, I_ℓ , specifies operations to be performed, including the initial operations $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, where said initial instructions have been transformed to a kernel, KE, of kernel instructions including the kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ in which execution commences with \bar{I}_0 one or more times, once for each iteration, \bar{i} , of the kernel, KE, where each said kernel instruction, \bar{I}_k , specifies kernel operations to be performed, including the kernel operations $\bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_N^{k,\ell}$, where each kernel operation, $\bar{O}_n^{k,\ell}$, provides address offsets, $ao_n^k(c)$, for use in connection with the kernel operations to be performed, and where instruction execution is measured by instruction periods, T , comprising,

instruction processing means for sequentially accessing said kernel instructions, \bar{I}_k , and corresponding kernel operations, $\bar{O}_n^{k,\ell}$, said instruction processing means accessing said instructions one or more times during one or more iterations, \bar{i} , of said instruction stream \bar{IS} ,

one or more processors, each processor for performing one or more of said kernel operations specified by said instructions, \bar{I}_k , each processor having one or more input ports for receiving input operands and output ports for providing result operands where a result operand from one

-109-

iteration of a loop is used as a source operand in a subsequent iteration of said loop, ,

a plurality of multiconnect memories, addressed by memory addresses, $a_n^k(c)(\bar{i})$, for connecting operands from and to said processors, said memories having input and output ports, said memories providing input operands on said memory output ports when addressed by said memory addresses,

processor-memory interconnection means for connecting output operands from processor output ports to memory input ports,

memory-processor interconnection means for connecting input operands from memory output ports to processor input ports,

invariant addressing means for addressing said memories during different iterations including a current iteration, \bar{i} , and a previous iteration, $(\bar{i}-1)$ including,

modifying means for forming a current pointer address, $mcp(\bar{i})$, from a previous pointer address $mcp(\bar{i}-1)$ with the operation $D^*[mcp(\bar{i}-1)]$ such that $mcp(\bar{i}) = D^*[mcp(\bar{i}-1)]$,

pointer register means for storing said pointer address, $mcp(\bar{i})$, for use in the \bar{i}^{th} -iteration,

address generation means for combining the pointer address, $mcp(\bar{i})$, with an address offset, $ao_n^k(c)$, to form said memory addresses, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration.

68. The system of Claim 67 wherein execution of the kernel is repeated a plurality of times including said $(\bar{i}-1)^{th}$ -iteration and said \bar{i}^{th} -iteration and wherein executions of the kernel causes the multiconnect pointer to change, wherein said kernel includes an instruction providing an address offset, ao_n^k , and wherein said modifying means modifies said previous pointer addresses, $mcp(\bar{i}-1)$,

-110-

in response to repetitions of said kernel to form said current pointer addresses, $mcp(\bar{i})$,

whereby said address generation means forms the memory address, $a_n^k(c)(\bar{i}-1)$, for said instruction, \bar{i}_k , for the $(\bar{i}-1)^{th}$ -iteration of said kernel as a function of said address offset ao_n^k and the previous pointer address $mcp(\bar{i}-1)$ and forms the memory address, $a_n^k(c)(\bar{i})$, for the i^{th} -iteration of said kernel as a function of said address offset $ao_n^k(c)$ and the current pointer address $mcp(\bar{i})$.

69. The system of Claim 68 wherein execution of the kernel occurs for iterations 1, 2, ..., $(\bar{i}-1)$, \bar{i} , ..., \bar{R} ,

said address generation means includes an adder for adding said current pointer address, $mcp(\bar{i})$, with offsets, ao_n^k , to form said memory addresses, $a_n^k(c)(\bar{i})$, and

said modifying means includes means to add one to said previous pointer address, $mcp(\bar{i}-1)$, for each iteration of said kernel to form said current pointer address equal to $mcp(1)$, $mcp(2)$, $mcp(3)$, ..., $mcp(\bar{i}-1)$, $mcp(\bar{i})$, ..., $mcp(\bar{Z})$,

whereby the memory addresses for said kernel instruction are determined for said kernel to form $ao_n^k(c)+mcp(1)$, $ao_n^k(c)+mcp(2)$, $ao_n^k(c)+mcp(3)$, ..., $ao_n^k(c) + mcp(\bar{i}-1)$, $ao_n^k(c) + mcp(i)$, ..., $ao_n^k(c)+mcp(X)$, respectively.

70. The system of Claim 67 wherein said instruction processing means includes,

a loop counter for storing a loop count, "lc" representing the number of iterations of said loop, LP, said loop count having a loop count range including a loop count end count,

an epilog counter for storing an epilog count, "esc" representing the ending iterations of said loop, KE, said epilog count having an epilog count range including an epilog end count,

counter control means for controlling the loop counter and epilog counter.

-111-

9. A computer system as set forth in Claim 5 wherein said instruction processing means includes,

means for storing an iteration value representing different iterations of said kernel, KE,

processor control means for enabling different ones of said processors as a function of said iteration value whereby different ones of the kernel operations from the same kernel instructions are executed during different iterations.

72. A computer system as set forth in Claim 71 herein said kernel, KE, and the kernel instructions, $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_K$, are iteratively executed during sequential executions of the kernel loop, KE, and wherein operations in each of the kernel instructions are selectively enabled during different iterations of the kernel loop, KE, under control of said processor control means.

73. A computer system as set forth in Claim 72 wherein said processor control means includes iteration control register means for storing control information for selectively enabling the operations, $\bar{O}_n^{k,\ell}$, during repetitive executions of the kernel, KE, and includes operation address means for addressing said iteration control register means.

74. A computer system as set forth in Claim 71 wherein said one or more processors include n processors, one for each of said n operations to be performed by said instructions, \bar{I}_k , said system further including iteration control means for enabling said processors selectively during different iterations.

75. The computer system of Claim 67 wherein said instruction processing means includes operation control means for selecting different ones of the operations to be

-112-

performed during each iteration \bar{i} of the kernel loop, KE and wherein said kernel loop includes a prolog, body and epilog.

76. The computer system of Claim 75 wherein said operation control means includes,

prolog control means for selecting an increasing number of operations to be performed for each successive iteration of \bar{i} during the prolog.

77. The computer system of Claim 75 wherein said operation control means includes,

body control means for selecting a constant number of operations to be performed during the body.

68. The computer system of Claim 67 wherein said counting means includes a loop counter for counting iterations \bar{i} over the count range R and includes an epilog stage counter for counting over the count range S-1 whereby said loop counter and said epilog stage counter together count over the count range \bar{R} equal to $R-(S-1)$.

79. The computer system of Claim 67 wherein each of said processors is associated with iteration control means to enable said processor in response to said instruction processing means during each selected iteration \bar{i} and wherein said instruction processing means includes means for selectively setting said iteration control means for enabling one or more operation in one or more of said processors during each iteration, \bar{i} .

80. The computer system of Claim 79 wherein said iteration control means is operative to control the operation in an enabled processor for each iteration \bar{i} of the kernel loop, KE.

-113-

81. The computer system of Claim 67 wherein said multiconnect memories includes means for storing results of each operation for each iteration in unique locations for each iteration.

82. The computer system of Claim 81 wherein said instruction processing means includes means for specifying the source for each operation of each iteration.

83. The computer system Claim 82 wherein said instruction processing means includes means for specifying the results for one or more previous iterations as a source for processing in the current iteration.

-114-

84. A computer system for performing one or more iterations of a loop of instructions, where the next instruction to be executed in one iteration of a loop is different in a subsequent iteration of said loop, including,

a processing unit having one or more processors for performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying for each instruction and each iteration of said loop of instructions operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer, for specifying next instructions and branch instructions,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit.

85. The system of Claim 84 wherein said instruction unit includes,

a loop counter for storing a loop count, "lc" representing the number of iterations of said loop, LP, said loop count having a loop count range including a loop count end count,

an epilog counter for storing an epilog count, "esc" representing the ending iterations of said loop, KE, said epilog count having an epilog count range including an epilog end count,

counter control means for controlling the loop counter and epilog counter.

-115-

86. A computer system as set forth in Claim 84 wherein said instruction unit includes,

means for storing an iteration value representing different iterations of said kernel, KE,

processor control means for enabling different ones of said processors as a function of said iteration value whereby different ones of the kernel operations from the same kernel instructions are executed during different iterations.

87. A computer system for performing one or more iterations of a loop of instructions, where a result operand from one iteration of a loop is used as a source operand in a subsequent iteration of said loop, including,

a processing unit having one or more processors for performing operations on input operands and providing output operands,

a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,

an instruction unit for specifying for each instruction and each iteration of said loop of instructions operations to be performed by said processing unit, for specifying source address offsets and destination address offsets relative to a modifiable pointer,

invariant addressing means for providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit, wherein said invariant addressing means stores the result operand at a first multiconnect address specified by a first address offset and a first value of said modifiable pointer during one iteration, and wherein said invariant addressing means accesses said result operand as a source operand from said multiconnect address specified by a combination of said

-116-

second address offset and a second value of said modifiable pointer during another iteration of said loop,

iteration control means for controlling which operations are active in each instruction during each iteration of said loop.

88. A computer system employing a horizontal architecture for use with an initial instruction stream, IS, where the initial instruction stream includes a loop, LP, of initial instructions including branch instructions, including initial instructions $I_0, I_1, I_2, \dots, I_\ell, \dots, I_{(L-1)}$ in which execution commences with I_0 one or more times, once for each iteration, i , of the loop, LP, where each said initial instruction, I_ℓ , specifies operations to be performed, including the initial operations $O_0^\ell, O_1^\ell, O_2^\ell, \dots, O_n^\ell, \dots, O_{(N-1)}^\ell$, where said initial instructions have been transformed to a kernel, KE, of kernel instructions including the kernel instructions $\bar{I}_0, \bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_{(K-1)}$ in which execution commences with \bar{I}_0 one or more times, once for each iteration, \bar{i} , of the kernel, KE, where each said kernel instruction, \bar{I}_k , specifies kernel operations to be performed, including the kernel operations $\bar{O}_1^{k,\ell}, \bar{O}_2^{k,\ell}, \dots, \bar{O}_n^{k,\ell}, \dots, \bar{O}_N^{k,\ell}$, where each kernel operation, $\bar{O}_n^{k,\ell}$, provides address offsets, $ao_n^k(c)$, for use in connection with the kernel operations to be performed, and where instruction execution is measured by instruction periods, T, comprising,

instruction processing means for sequentially accessing said kernel instructions, \bar{I}_k , and corresponding kernel operations, $\bar{O}_n^{k,\ell}$, said instruction processing means accessing said instructions one or more times during one or more iterations, \bar{i} , of said instruction stream \bar{IS} ,

one or more processors, each processor for performing one or more of said kernel operations specified by said instructions, \bar{I}_k , each processor having one or more input ports for receiving input operands and output ports for

-117-

providing result operands where a result operand from one iteration of a loop is used as a source operand in a subsequent iteration of said loop, ,

a plurality of multiconnect memories, addressed by memory addresses, $a_n^k(c)(\bar{i})$, for connecting operands from and to said processors, said memories having input and output ports, said memories providing input operands on said memory output ports when addressed by said memory addresses,

processor-memory interconnection means for connecting output operands from processor output ports to memory input ports,

memory-processor interconnection means for connecting input operands from memory output ports to processor input ports,

invariant addressing means for addressing said memories during different iterations including a current iteration, \bar{i} , and a previous iteration, $(\bar{i}-1)$ including,

modifying means for forming a current pointer address, $mcp(\bar{i})$, from a previous pointer address $mcp(\bar{i}-1)$ with the operation $D^*[mcp(\bar{i}-1)]$ such that $mcp(\bar{i})=D^*[mcp(\bar{i}-1)]$,

pointer register means for storing said pointer address, $mcp(\bar{i})$, for use in the \bar{i}^{th} -iteration,

address generation means for combining the pointer address, $mcp(\bar{i})$, with an address offset, $ao_n^k(c)$, to form said memory addresses, $a_n^k(c)(\bar{i})$, for the \bar{i}^{th} -iteration.

89. The system of Claim 88 wherein execution of the kernel is repeated a plurality of times including said $(\bar{i}-1)^{th}$ -iteration and said \bar{i}^{th} -iteration and wherein executions of the kernel causes the multiconnect pointer to change, wherein said kernel includes an instruction providing an address offset, ao_n^k , and wherein said modifying

-118-

means modifies said previous pointer addresses, $mcp(\bar{i}-1)$, in response to repetitions of said kernel to form said current pointer addresses, $mcp(\bar{i})$,

whereby said address generation means forms the memory address, $a_n^k(c)(\bar{i}-1)$, for said instruction, \bar{i}_k , for the $(\bar{i}-1)^{th}$ -iteration of said kernel as a function of said address offset ao_n^k and the previous pointer address $mcp(\bar{i}-1)$ and forms the memory address, $a_n^k(c)(\bar{i})$, for the i^{th} -iteration of said kernel as a function of said address offset $ao_n^k(c)$ and the current pointer address $mcp(\bar{i})$.

90. The system of Claim 89 wherein execution of the kernel occurs for iterations 1, 2, ..., $(\bar{i}-1)$, \bar{i} , ..., \bar{R} ,

said address generation means includes an adder for adding said current pointer address, $mcp(\bar{i})$, with offsets, ao_n^k , to form said memory addresses, $a_n^k(c)(\bar{i})$, and

said modifying means includes means to add one to said previous pointer address, $mcp(\bar{i}-1)$, for each iteration of said kernel to form said current pointer address equal to $mcp(1)$, $mcp(2)$, $mcp(3)$, ..., $mcp(\bar{i}-1)$, $mcp(\bar{i})$, ..., $mcp(\bar{Z})$,

whereby the memory addresses for said kernel instruction are determined for said kernel to form $ao_n^k(c)+mcp(1)$, $ao_n^k(c)+mcp(2)$, $ao_n^k(c)+mcp(3)$, ..., $ao_n^k(c) + mcp(\bar{i}-1)$, $ao_n^k(c) + mcp(i)$, ..., $ao_n^k(c)+mcp(X)$, respectively.

91. The system of Claim 88 wherein said instruction processing means includes,

a loop counter for storing a loop count, "lc" representing the number of iterations of said loop, LP, said loop count having a loop count range including a loop count end count,

an epilog counter for storing an epilog count, "esc" representing the ending iterations of said loop, KE, said epilog count having an epilog count range including an epilog end count,

-119-

counter control means for controlling the loop counter and epilog counter.

92. A computer system as set forth in Claim 88 wherein said instruction processing means includes,

means for storing an iteration value representing different iterations of said kernel, KE,

processor control means for enabling different ones of said processors as a function of said iteration value whereby different ones of the kernel operations from the same kernel instructions are executed during different iterations.

93. A computer system as set forth in Claim 92 herein said kernel, KE, and the kernel instructions, $\bar{I}_1, \bar{I}_2, \dots, \bar{I}_k, \dots, \bar{I}_K$, are iteratively executed during sequential executions of the kernel loop, KE, and wherein operations in each of the kernel instructions are selectively enabled during different iterations of the kernel loop, KE, under control of said processor control means.

94. A computer system as set forth in Claim 93 wherein said processor control means includes iteration control register means for storing control information for selectively enabling the operations, $\bar{O}_n^{k,l}$ during repetitive executions of the kernel, KE, and includes operation address means for addressing said iteration control register means.

95. A computer system as set forth in Claim 92 wherein said one or more processors include n processors, one for each of said n operations to be performed by said instructions, \bar{I}_k , said system further including iteration control means for enabling said processors selectively during different iterations.

-120-

96. The computer system of Claim 88 wherein said instruction processing means includes operation control means for selecting different ones of the operations to be performed during each iteration \bar{i} of the kernel loop, KE and wherein said kernel loop includes a prolog, body and epilog.

97. The computer system of Claim 96 wherein said operation control means includes,

prolog control means for selecting an increasing number of operations to be performed for each successive iteration of \bar{i} during the prolog.

98. The computer system of Claim 96 wherein said operation control means includes,

body control means for selecting a constant number of operations to be performed during the body.

99. The computer system of Claim 88 wherein said counting means includes a loop counter for counting iterations \bar{i} over the count range R and includes an epilog stage counter for counting over the count range S-1 whereby said loop counter and said epilog stage counter together count over the count range \bar{R} equal to $R-(S-1)$.

100. The computer system of Claim 88 wherein each of said processors is associated with iteration control means to enable said processor in response to said instruction processing means during each selected iteration \bar{i} and wherein said instruction processing means includes means for selectively setting said iteration control means for enabling one or more operation in one or more of said processors during each iteration, \bar{i} .

101. The computer system of Claim 100 wherein said iteration control means is operative to control the operation

-121-

in an enabled processor for each iteration \bar{i} of the kernel loop, KE.

102. The computer system of Claim 88 wherein said multiconnect memories includes means for storing results of each operation for each iteration in unique locations for each iteration.

103. The computer system of Claim 102 wherein said instruction processing means includes means for specifying the source for each operation of each iteration.

104. The computer system Claim 104 wherein said instruction processing means includes means for specifying the results for one or more previous iterations as a source for processing in the current iteration.

105. A computer system including,
a processing unit having a plurality of processors for performing operations on input operands and providing output operands,
a multiconnect unit for storing operands at addressable locations and for providing said input operands from source addresses and for storing said output operands with destination addresses,
an instruction unit having a plurality of locations for specifying an operation to be performed by said processors, each for specifying source address offsets and destination address offsets relative to a modifiable pointer,
connector means for connecting each of said locations to said corresponding processors,

-122-

invariant addressing means providing said modifiable pointer and for combining said address offsets to form said source addresses and said destination addresses in said multiconnect unit.

106. The computer system of Claim 105, wherein said connection means connects said locations to said processors concurrently and in parallel.

107. The system of Claim 105, wherein said connection means connects said locations to said processors one at a time.

108. The computer system of Claim 105, wherein said connection means includes first connection means for connecting said locations to said processors concurrently and in parallel and includes second connection means for connecting said locations to said processors one at a time and wherein said instruction means includes means for selecting between said first and second connection means.




FIG. -1

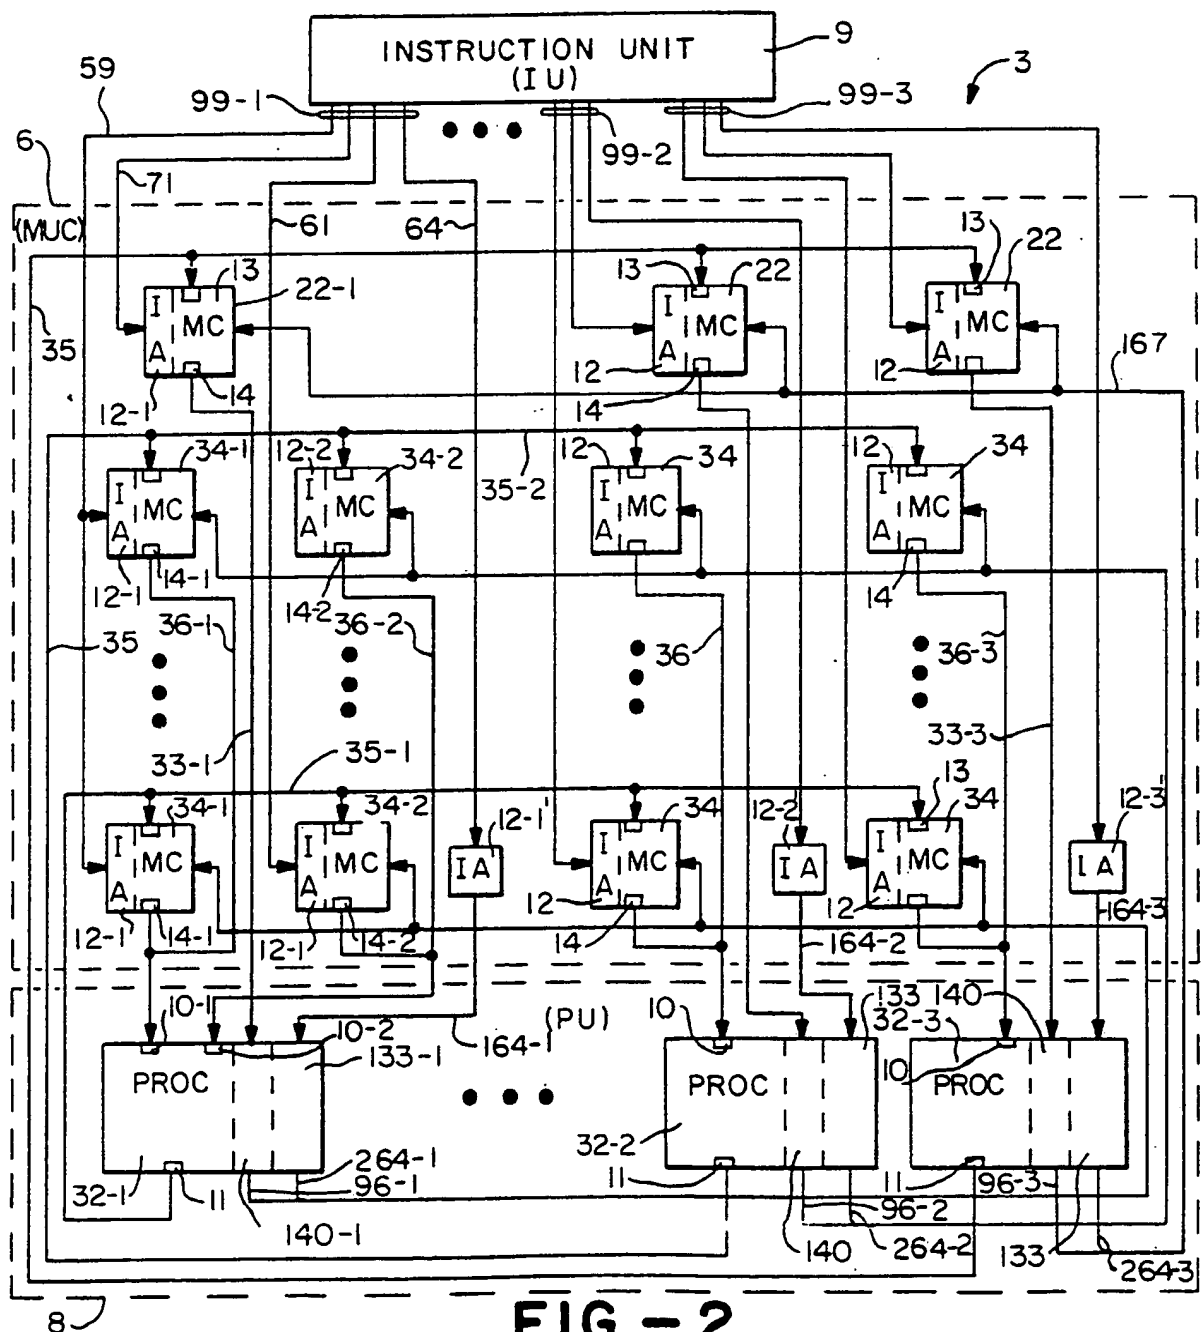


FIG. -2

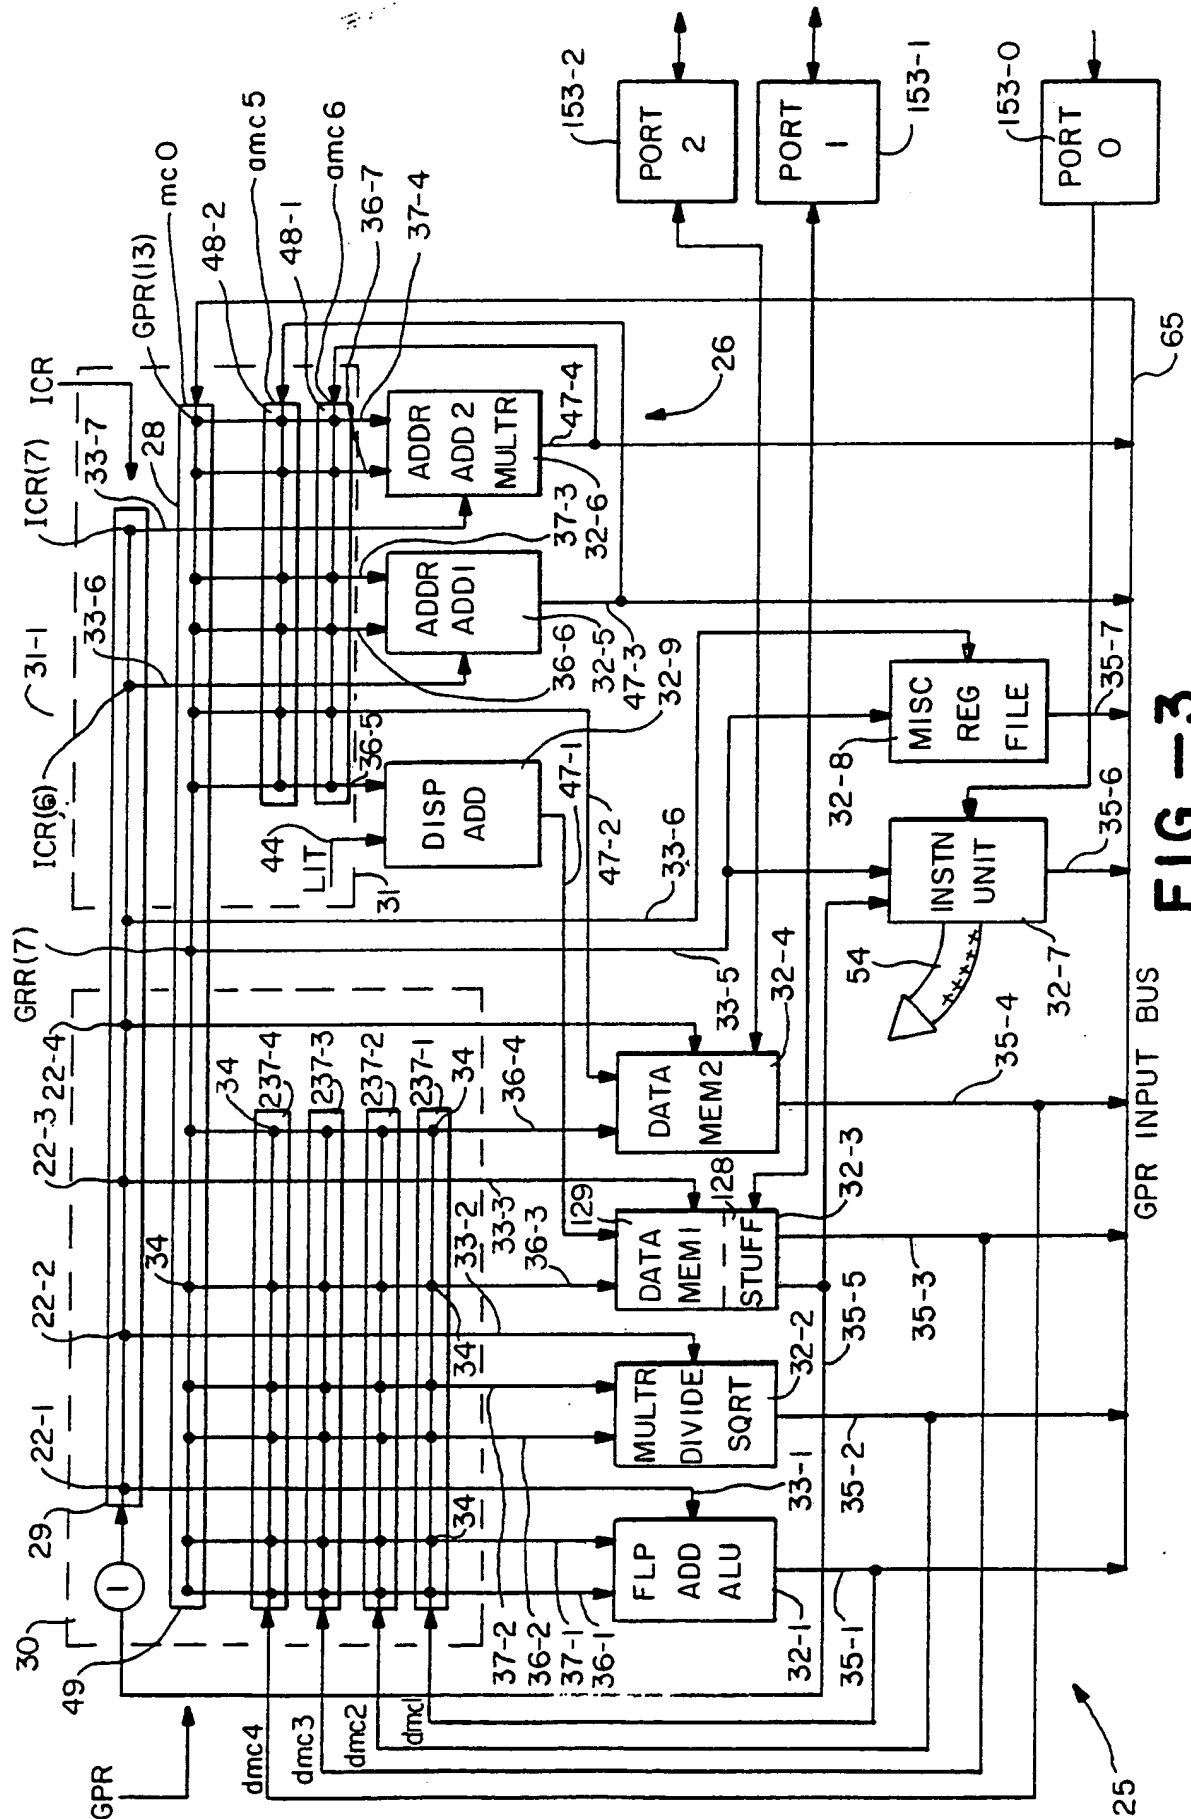
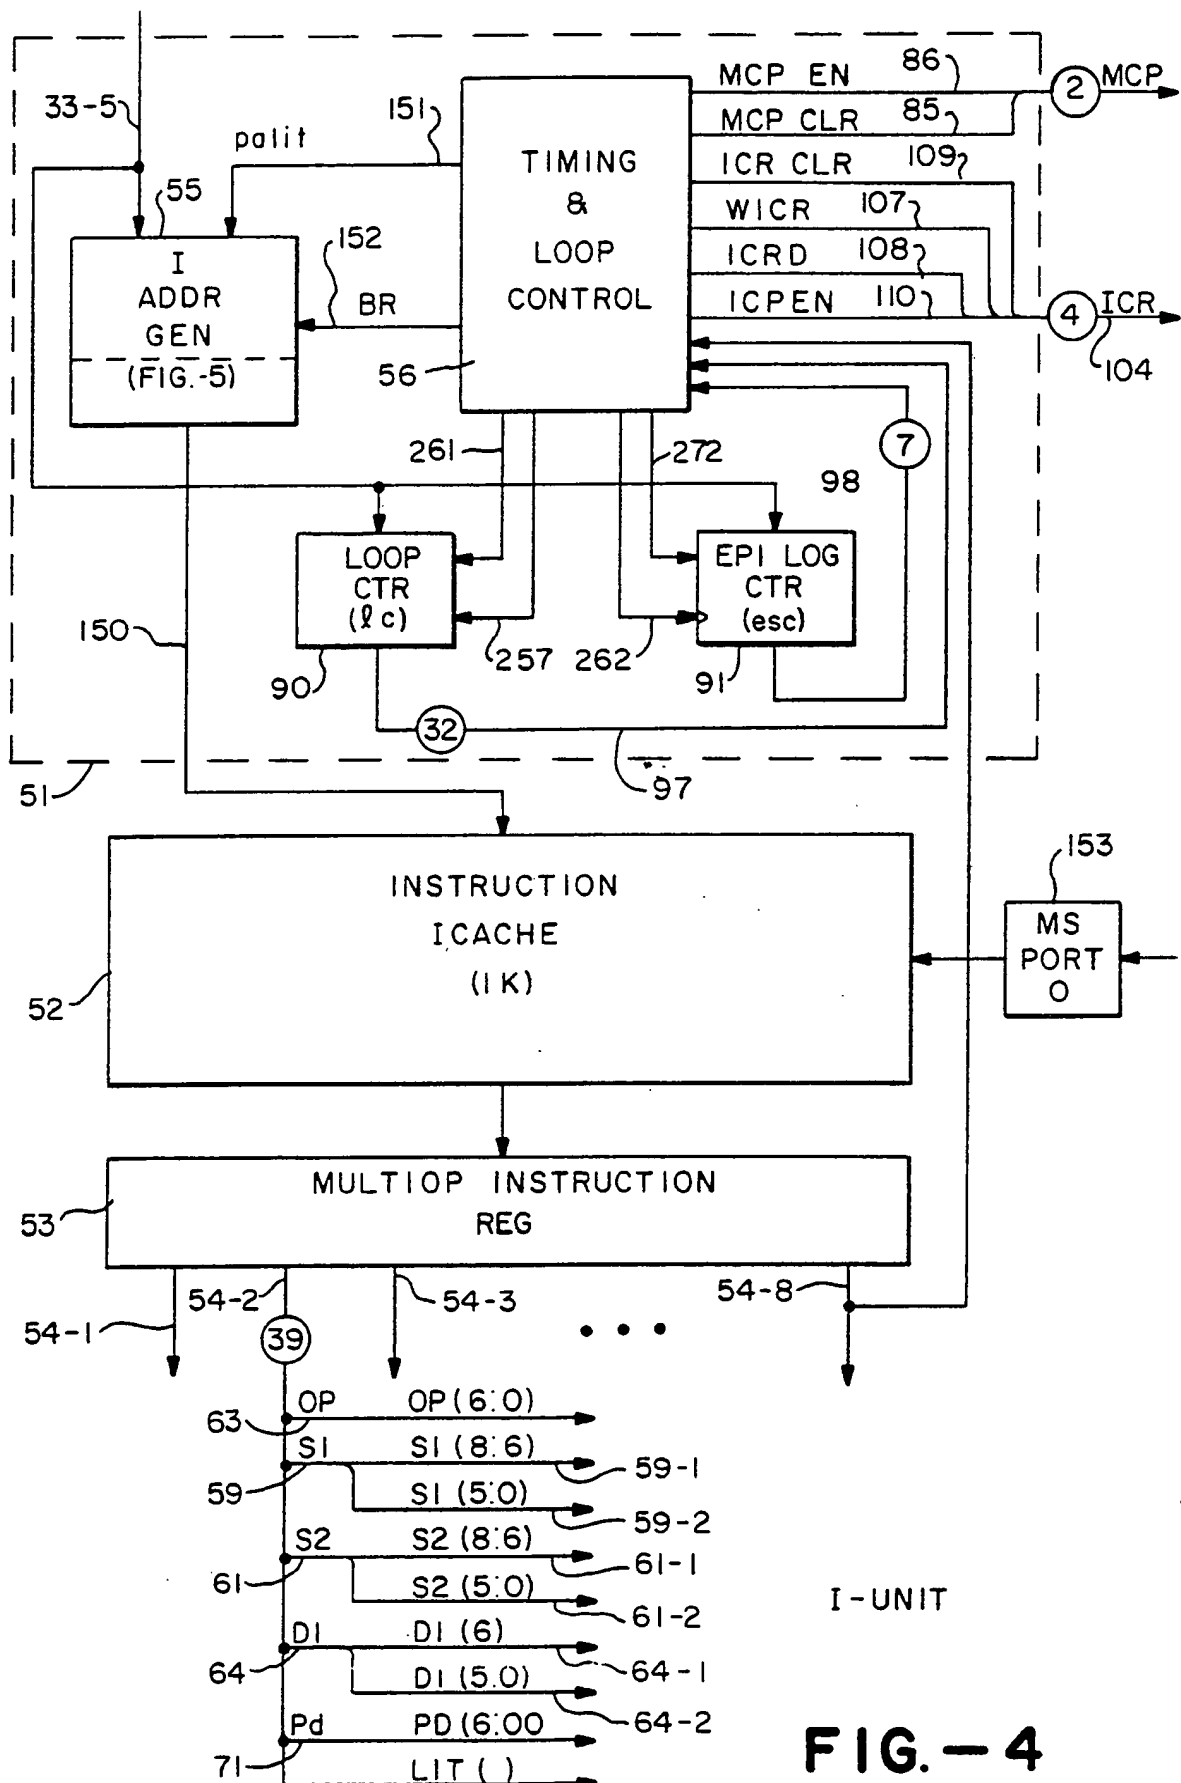


FIG.-3






FIG. -5

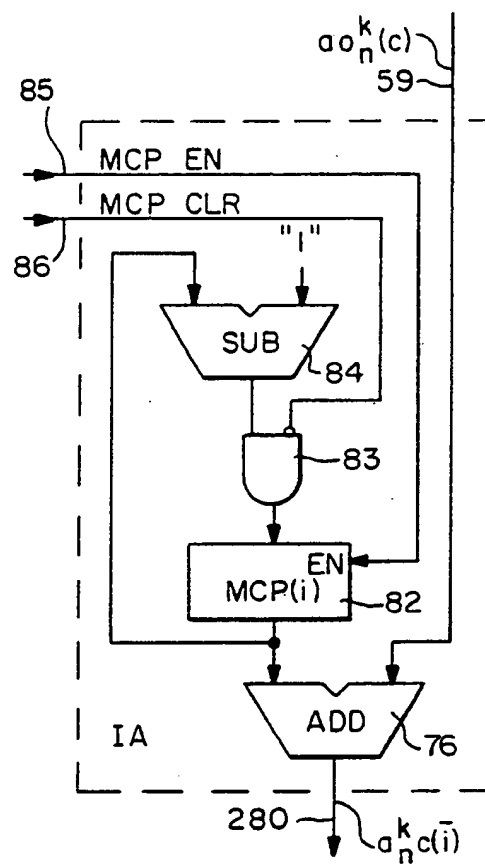


FIG.-6


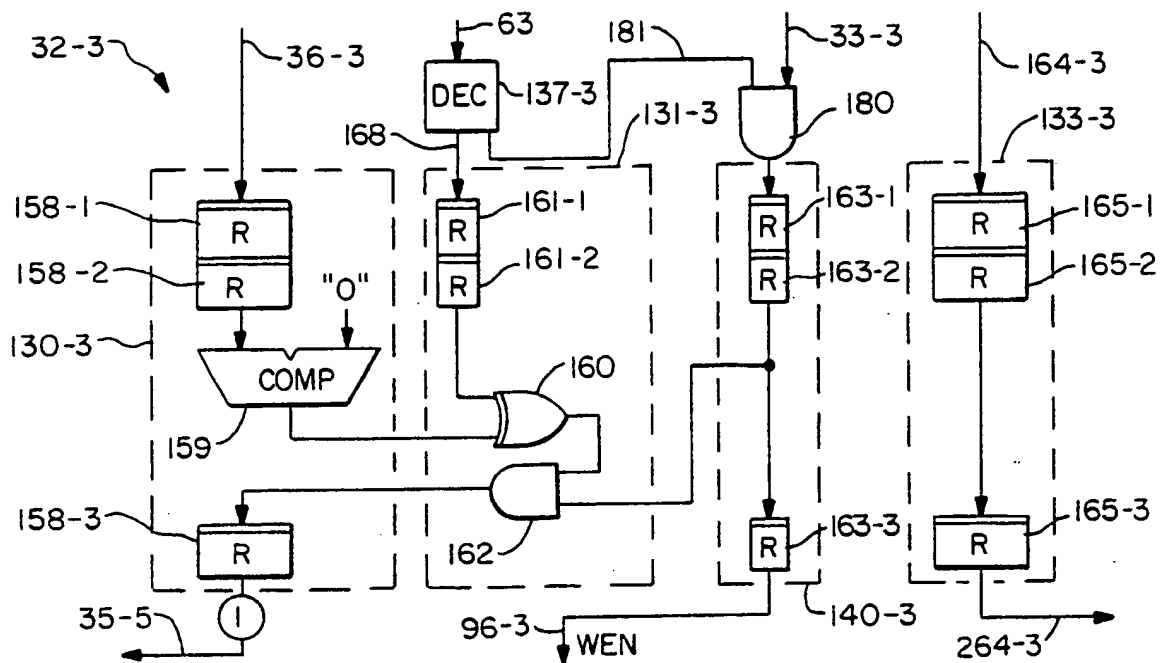



FIG. -7

**FIG. - 8****FIG. - 9**

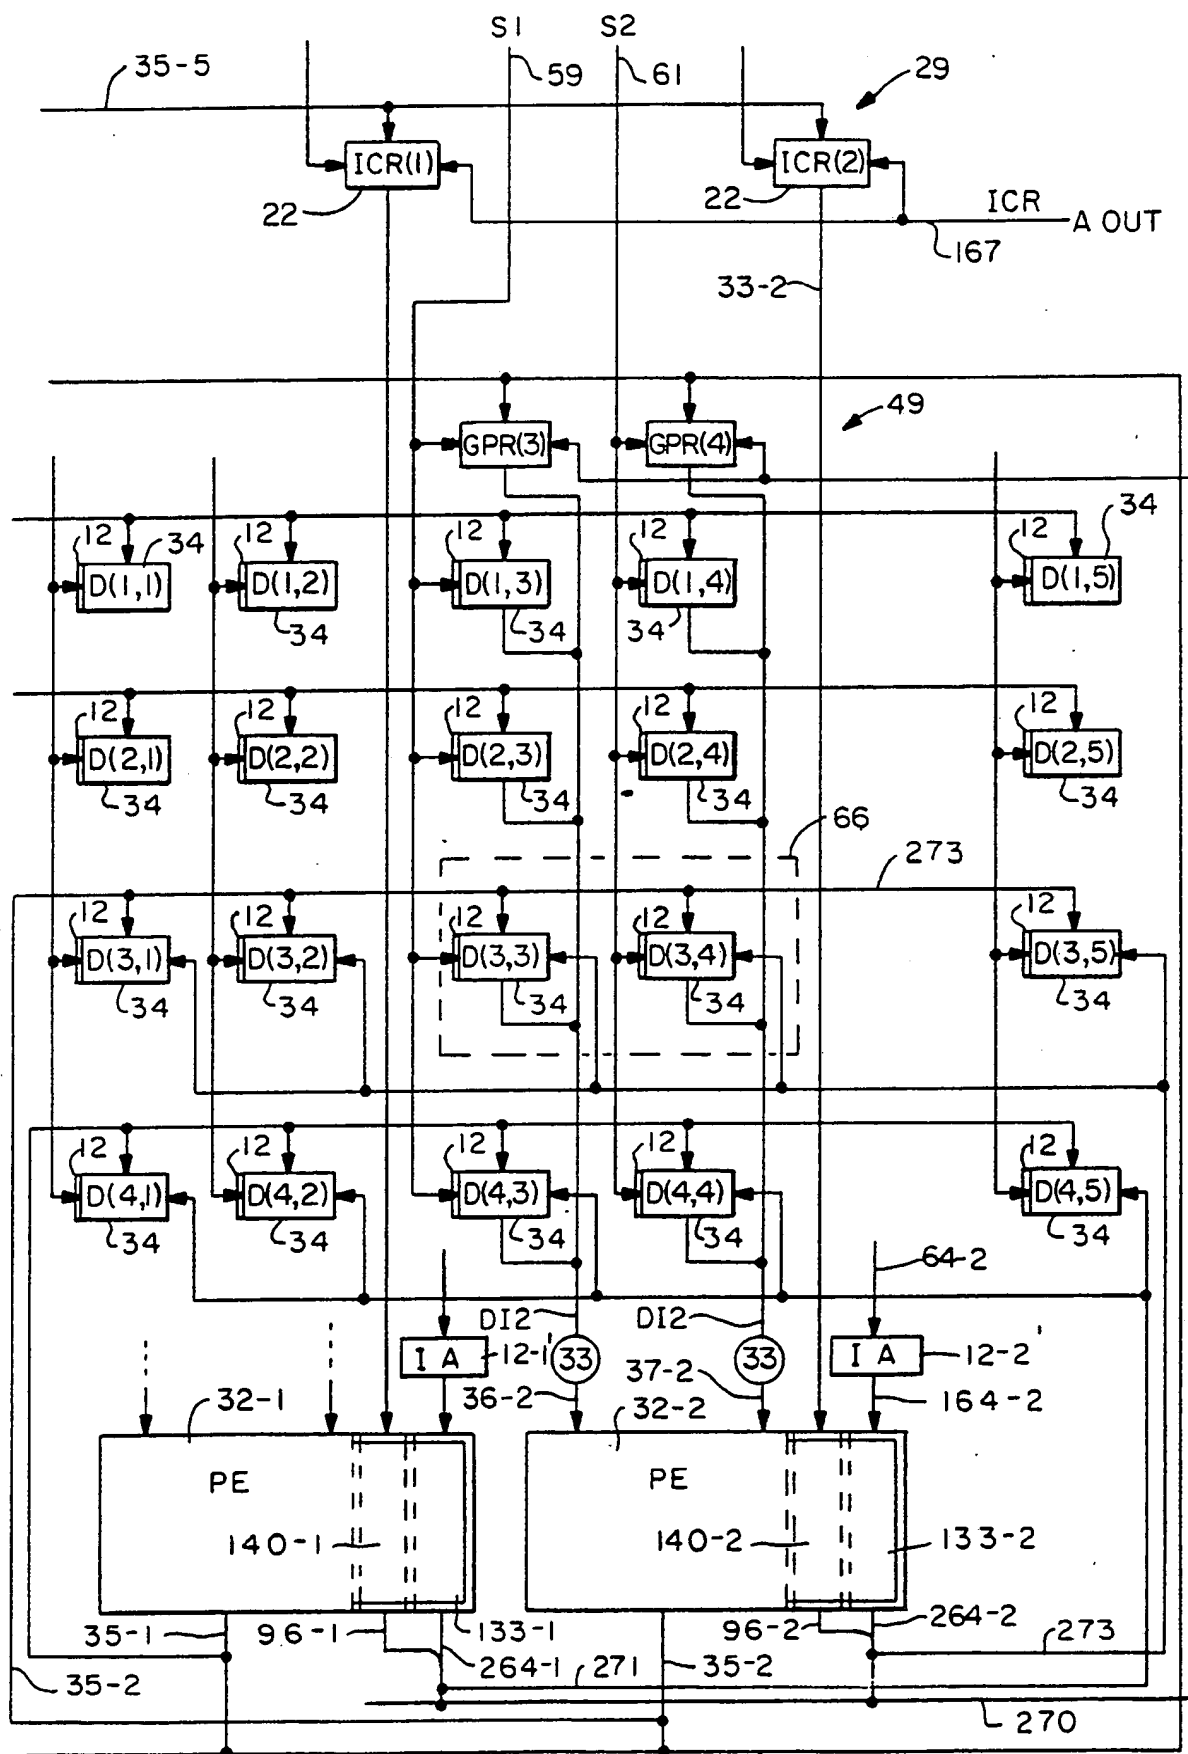


FIG. -10

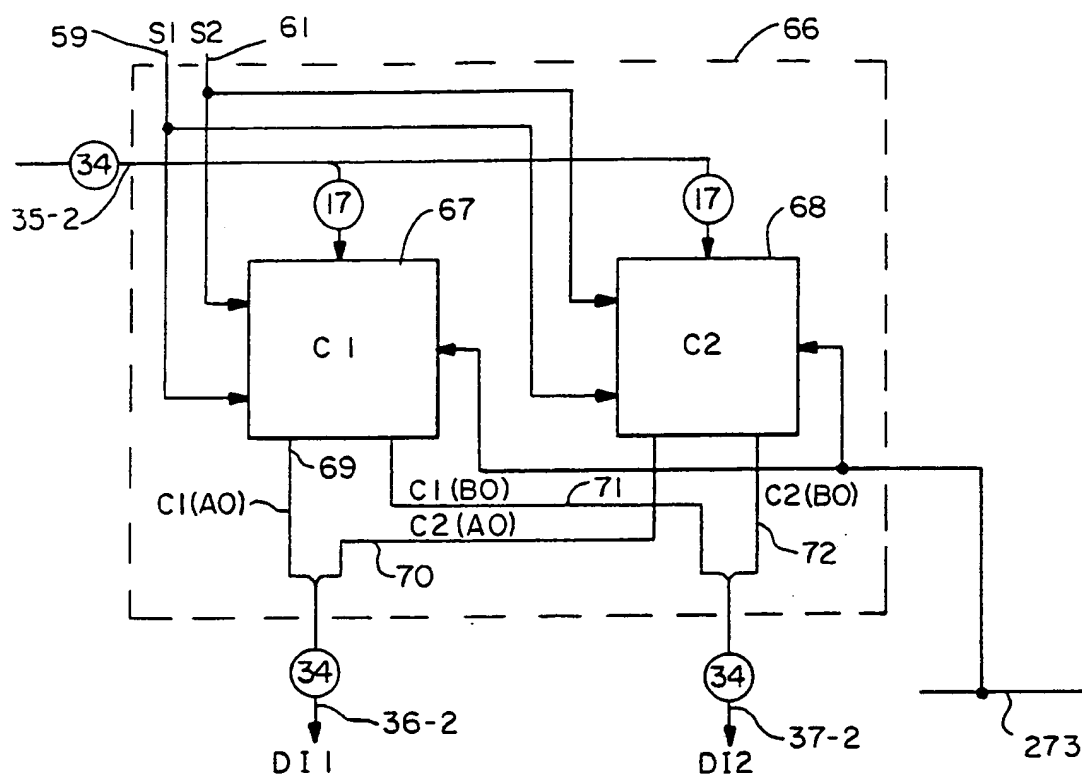


FIG. - II




FIG. -15

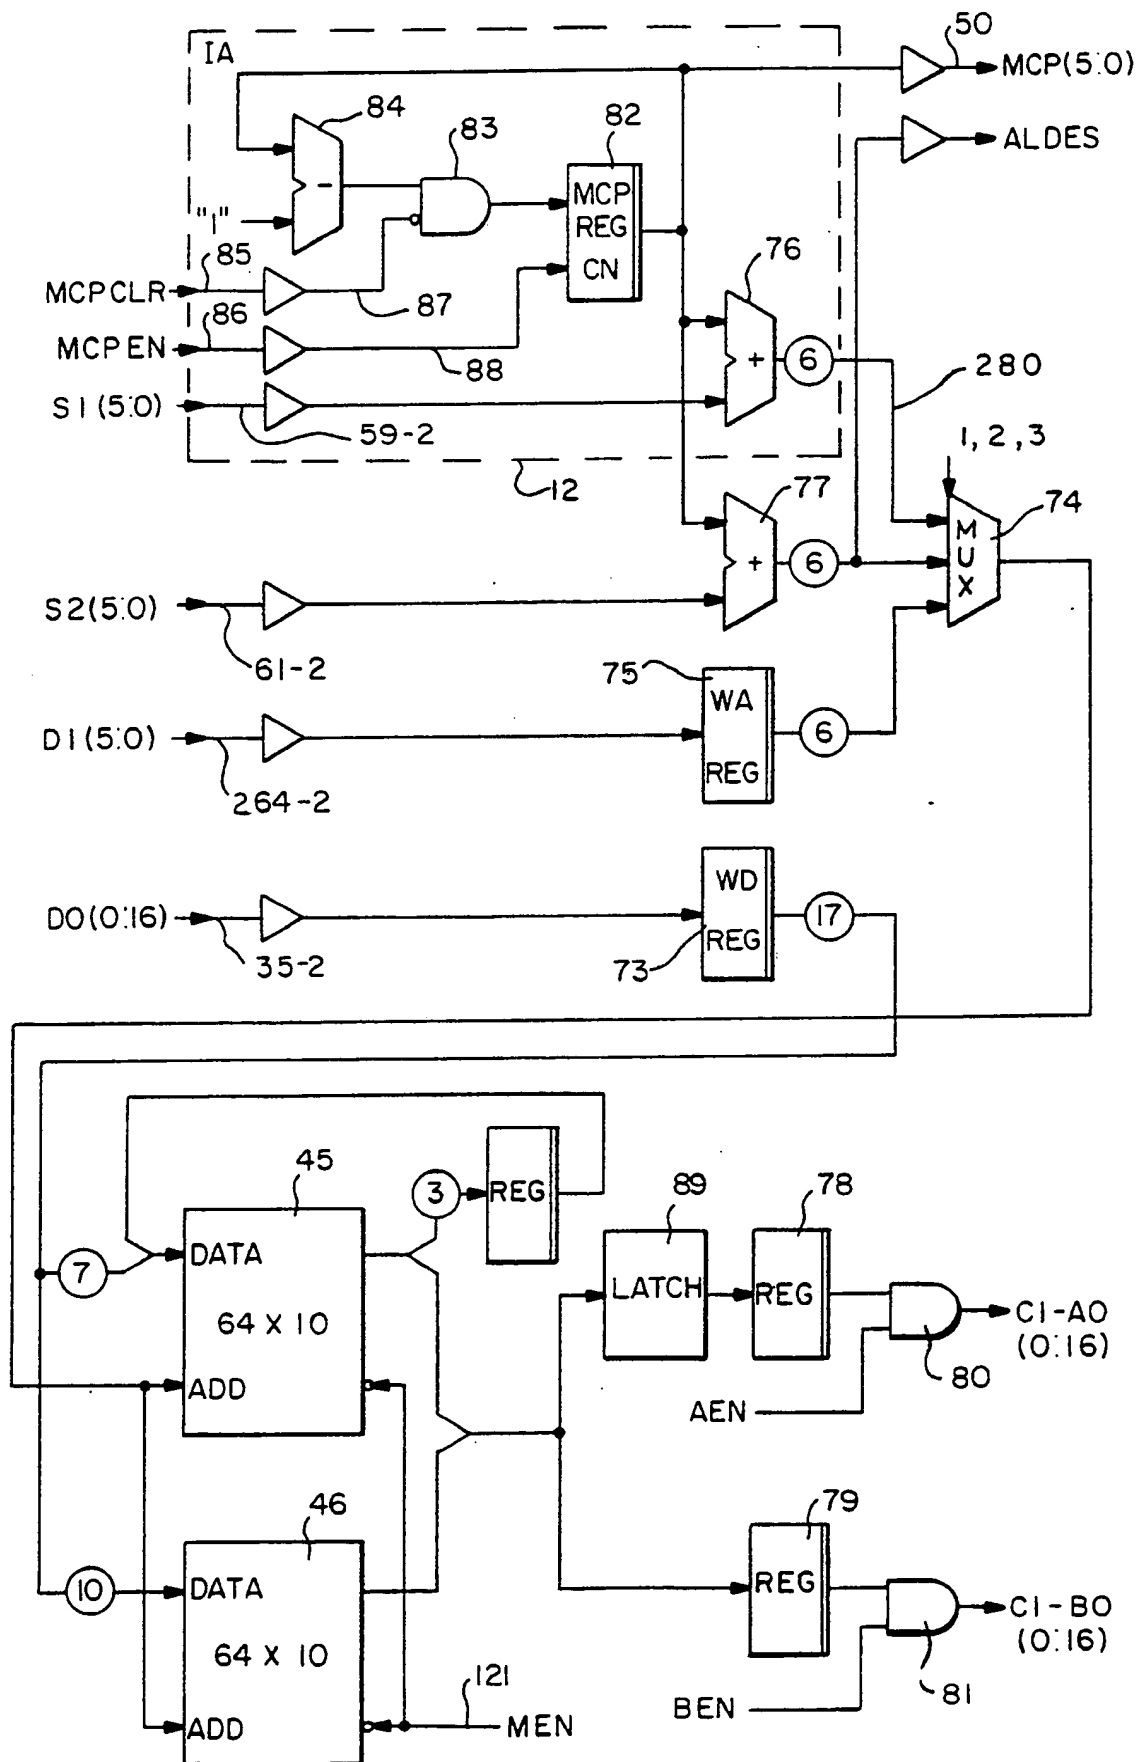


FIG. - 12




FIG. -13




FIG. -14

INTERNATIONAL SEARCH REPORT

International Application No. PCT/US88/01413

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ⁶ According to International Patent Classification (IPC) or to both National Classification and IPC IPC(4): G06F 9/38, 13/02 U.S. Cl. 364/200																	
II. FIELDS SEARCHED <div style="text-align: center; border-top: 1px solid black; border-bottom: 1px solid black; margin: 5px 0;">Minimum Documentation Searched ⁷</div> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 25%; border-bottom: 1px solid black; padding: 5px;">Classification System</td> <td style="border-bottom: 1px solid black; padding: 5px;">Classification Symbols</td> </tr> <tr> <td style="padding: 5px;">U.S.</td> <td style="padding: 5px;">364/200, 900</td> </tr> </table> <div style="text-align: center; border-top: 1px solid black; border-bottom: 1px solid black; margin: 5px 0;">Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸</div>			Classification System	Classification Symbols	U.S.	364/200, 900											
Classification System	Classification Symbols																
U.S.	364/200, 900																
III. DOCUMENTS CONSIDERED TO BE RELEVANT ⁹ <table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 10%; border-bottom: 1px solid black; padding: 5px;">Category [*]</th> <th style="width: 60%; border-bottom: 1px solid black; padding: 5px;">Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²</th> <th style="width: 30%; border-bottom: 1px solid black; padding: 5px;">Relevant to Claim No. ¹³</th> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">X Y A</td> <td style="padding: 5px;">US, A, 4,553,203 (RAU) 12 November 1985, see entire document.</td> <td style="padding: 5px;">1, 61, 63, 66, 64, 87, 105-108 2, 5-6, 7-17, 19-20, 62, 65, 67-69, 86 18, 21-27, 23, 33-60, 72-83, 88-90, 94, 95-98, 100-104</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">Y A</td> <td style="padding: 5px;">US, A, 4,292,667 (CATILLER) 29 September 1985, see column 14, lines 25-46.</td> <td style="padding: 5px;">3-4, 64, 70, 85, 91 32, 71, 92, 99</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A, P</td> <td style="padding: 5px;">US, A, 4,740,894 (LYON) 26 April 1988, see entire document.</td> <td style="padding: 5px;">1-108</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A</td> <td style="padding: 5px;">US, A, 4,455,938 (PARKER) 3 December 1985, see entire document.</td> <td style="padding: 5px;">3-4, 32, 70-71, 85, 91-92, 99</td> </tr> </table>			Category [*]	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³	X Y A	US, A, 4,553,203 (RAU) 12 November 1985, see entire document.	1, 61, 63, 66, 64, 87, 105-108 2, 5-6, 7-17, 19-20, 62, 65, 67-69, 86 18, 21-27, 23, 33-60, 72-83, 88-90, 94, 95-98, 100-104	Y A	US, A, 4,292,667 (CATILLER) 29 September 1985, see column 14, lines 25-46.	3-4, 64, 70, 85, 91 32, 71, 92, 99	A, P	US, A, 4,740,894 (LYON) 26 April 1988, see entire document.	1-108	A	US, A, 4,455,938 (PARKER) 3 December 1985, see entire document.	3-4, 32, 70-71, 85, 91-92, 99
Category [*]	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³															
X Y A	US, A, 4,553,203 (RAU) 12 November 1985, see entire document.	1, 61, 63, 66, 64, 87, 105-108 2, 5-6, 7-17, 19-20, 62, 65, 67-69, 86 18, 21-27, 23, 33-60, 72-83, 88-90, 94, 95-98, 100-104															
Y A	US, A, 4,292,667 (CATILLER) 29 September 1985, see column 14, lines 25-46.	3-4, 64, 70, 85, 91 32, 71, 92, 99															
A, P	US, A, 4,740,894 (LYON) 26 April 1988, see entire document.	1-108															
A	US, A, 4,455,938 (PARKER) 3 December 1985, see entire document.	3-4, 32, 70-71, 85, 91-92, 99															
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>[*] Special categories of cited documents: ¹⁰</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </div> <div style="width: 45%;"> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&" document member of the same patent family</p> </div> </div>																	
IV. CERTIFICATION <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; border-bottom: 1px solid black; padding: 5px;">Date of the Actual Completion of the International Search</td> <td style="width: 50%; border-bottom: 1px solid black; padding: 5px;">Date of Mailing of this International Search Report</td> </tr> <tr> <td style="padding: 5px;">23 June 1988</td> <td style="padding: 5px; text-align: center;">12 JUL 1988</td> </tr> <tr> <td style="border-bottom: 1px solid black; padding: 5px;">International Searching Authority</td> <td style="border-bottom: 1px solid black; padding: 5px;">Signature of Authorized Officer</td> </tr> <tr> <td style="padding: 5px;">ISA/US</td> <td style="padding: 5px; text-align: center;"> Ayni Mohamed </td> </tr> </table>			Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report	23 June 1988	12 JUL 1988	International Searching Authority	Signature of Authorized Officer	ISA/US	 Ayni Mohamed							
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report																
23 June 1988	12 JUL 1988																
International Searching Authority	Signature of Authorized Officer																
ISA/US	 Ayni Mohamed																

III. DOCUMENTS CONSIDERED TO BE RELEVANT (CONTINUED FROM THE SECOND SHEET)

Category *	Citation of Document, with indication, where appropriate, of the relevant passages	Relevant to Claim No
A	US, A, 4,310,879 (PANDEYA) 12 January 1982, see entire document.	1-108